AFRL-RV-PS-TR-2014-0075

AFRL-RV-PS-TR-2024-0075

ADAPTING PIPELINE ARCHITECTURES TO TRACK DEVELOPING AFTERSHOCK SEQUENCES AND RECURRENT EXPLOSIONS

Tormod Kværna, et al.

NORSAR
PO Box 53
N-2027 Kjeller, Norway

14 February 2014

Final Report

AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776

DTIC COPY

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 377 ABW Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RV-PS-TR-2014-0075 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//SIGNED//	//SIGNED//
Robert Raistrick	Edward J. Masterson, Colonel, USAF
Project Manager, AFRL/RVBYE	Chief, Battlespace Environment Division

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
14-02-2014	Final Report	11 Jan 2011 to 10 Jan 2014

4. TITLE AND SUBTITLE

Adapting Pipeline Architectures to Track Developing Aftershock Sequences and Recurrent Explosions

5a. CONTRACT NUMBER
FA9453-11-C-0230

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
62601F

6. AUTHOR(S)
T. Kværna, D. B. Harris[1], S. J. Gibbons, and D. Dodge[2]

5d. PROJECT NUMBER
1010

5e. TASK NUMBER
PPM00009357

5f. WORK UNIT NUMBER
EF004099

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
NORSAR
PO Box 53
N-2027 Kjeller, Norway

[1]Deschutes Signal Processing, LLC
81211 E Wapinitia Road
Maupin, Oregon 97037

[2] Lawrence Livermore National Laboratory
7000 East Avenue
Livermore, CA 94550

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory
Space Vehicles Directorate
3550 Aberdeen Avenue SE
Kirtland AFB, NM 87117-5776

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RVBYE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-RV-PS-TR-2014-0075

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited. (377ABW-2014-0329 dtd 08 May 2014)

13. SUPPLEMENTARY NOTES

14. ABSTRACT
A detection framework has been developed which augments a routine processing pipeline with mechanisms for spawning and administering so-called pattern detectors to identify and organize repeating waveforms discovered in multichannel seismic data streams. The framework has been tested and evaluated on a variety of different test cases from mining blasts in Central Asia to moderate and large earthquake aftershock sequences. The framework performs exceptionally in identifying repeating mining blasts, with signals from large clusters of events belonging exclusively to known quarries being detected and grouped fully automatically. For the moderate aftershock sequence scenario, the framework provided an event catalog down to a significantly lower completeness magnitude than could be generated with the available analyst resources. The large earthquake scenarios are far more challenging given the exceptionally large source regions covered by the aftershocks. The formation of high-rank subspace detectors can sweep up large proportions of the seismicity although not necessarily in a way that reduces significantly the burden of analyst interpretation. We suggest the use of a subspace-measure of waveform similarity that may perform better than the classical correlation coefficient for forming and evaluating clusters in such cases. Submerging a signal from the DPRK nuclear test into data from the same array from an extensive aftershock sequence, we demonstrate that this important signal is not screened out by the pattern detector components of the framework. We propose empirical matched field processing (EMFP) on single array streams as a sensitive primary detector for generating triggers at, and only at, the times of arrivals from events in the region of interest.

15. SUBJECT TERMS
Aftershock Sequences, Recurrent Explosions, Seismic Arrays, Subspace Detectors, Empirical Matched Field Processing

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Robert Raistrick
Unclassified	Unclassified	Unclassified	Unlimited	110	**19b. TELEPHONE NUMBER** *(include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

This page is intentionally left blank.

Table of Contents

Summary 1
Acknowledgements 2
Project Objectives 2
1 Introduction 4
2 The Detection Framework 8
2.1 Framework Architecture 8
2.2 Research Environment – Supplementary Software 15
3 Evaluation of Framework Performance 19
3.1 The 2005 Kashmir sequence 20
3.2 The 2012 Sumatra sequence 31
3.3 The 2008 Storfjorden sequence 38
3.4 Repeating mining explosions in Western Kazakhstan 44
4 Special Topics of Investigation 47
4.1 Case scenario of detection of DPRK signals within the Kashmir aftershock sequence 48
4.2 Subspace Clustering of the Tohoku Sequence 55
4.3 Investigation of single-phase empirical matched field processing as primary detectors 65
4.4 Necessity of Screening Criteria for Template Definition 77
4.5 Mitigation of the effects of interfering signals in pattern detector templates 79
5 Conclusions 85
References 87
Appendix - Relocation of events in the Kashmir aftershock sequence 89
Symbols, Abbreviations, and Acronyms 95

List of Figures

Fig. 2.1.1 The high-level block diagram of the detection framework as proposed and actually built. The boxes in yellow indicate functions approximately shared with a conventional pipeline detection processor. Those in green represent entirely new functions. Note that some detectors may process data from more than one array, possibly coherently..... 8
Fig. 2.2.1 Simplified schematic of the research workflow surrounding execution of the framework. As the framework has evolved to support a number of research objectives at 3 institutions it has become more flexible and complex, with a large number of configuration parameters. Its operation can store the triggers and detections associated with thousands of detectors. Four tools have been developed to manage the workflow..... 15
Fig. 2.2.2 Subdirectories and parameter text files created by ConfigCreator to support the execution of the framework. The framework actually is run using the command-line script framework_ runner..... 16
Fig. 2.2.3 Screenshot of the diagnostic GUI optionally displayed as the framework is running. . 17
Fig. 2.2.4 Screenshot of the Builder GUI used to review waveforms of events detected by a particular detector. The detector is selected from the list displayed in the left panel and the event waveforms appear in the right panel..... 18
Fig. 3.1.1 Map showing the IDC Reviewed Event Bulletin (REB) locations of the 2005 Kashmir earthquake aftershock sequence (grey symbols) together with KKAR array (red filled circle) and other stations in the region recording the events (triangles)..... 20
Fig. 3.1.2 Recordings of three different Kashmir aftershocks at the KK01 vertical component sensor of the KKAR array. The data are bandpass filtered between 1 and 3 Hz. 21
Fig. 3.1.3 Left: Waveforms at the KK01_SHZ sensor of the 18 events of the largest event cluster. Right: Locations of the 18 events (green filled circles) shown within the locations of the remaining events in the reference database (black filled circles). 22
Fig. 3.1.4 Left: Waveforms at the KK01_SHZ sensor of the 7 events of the second largest event cluster. Right: Locations of the 7 events (green filled circles) shown within the locations of the remaining events in the reference database (black filled circles)..... 23
Fig. 3.1.5 The upper panel displays the mutual KKAR cross-correlations among the 786 events in the reference database ordered after cluster analysis. The scale ranges linearly between 0 and 1. The lower panel shows a zoom-in on the first 260 events within the red box of the upper panel..... 24

Fig. 3.1.6 Analysis of waveform segments surrounding primary detections in the “Builder” program. Manual inspection identifies instances of templates that, for example, include foreign signals which would degrade the performance of a subspace detector spanning a representative set of events. Examples of these interfering signals are highlighted in red..... 26
Fig. 3.1.7 Display from the Builder program with the “clean” events used for creating the high-rank subspace detector. The shaded area indicates the time window used (approximately 140 seconds). The data were bandpass filtered between 1 and 3 Hz. 26
Fig. 3.1.8 Locations of the 551 events in the reference database detected by the high-rank subspace detector on the KKAR array in the time period 2005-270 to 2005-291 (green symbols) and events in the reference database not detected by the subspace detector (black symbols). 27
Fig. 3.1.9 KK01 SHZ waveforms of the last 50 of the 222 additional events found by the high-rank subspace detector..... 28
Fig. 3.1.10 KK01 SHZ waveforms of the 50 latest aftershock reference events missed by the high-rank subspace detector. 28
Fig. 3.1.11 Selected KK01 SHZ waveforms of missed events with overlapping aftershocks within the 150 second time window. 30
Fig. 3.2.1 Left: The 1222 events from the REB in the source region for 44 days following the main shock. (Right) Events per day for this dataset. Note that by day ten only a few tens of events are being produced in the source region. 31
Fig. 3.2.2 (White) Events from bulletin. (Blue) Correlation Detections. (Red) Subspace Detections..... 33
Fig. 3.2.3 Seismograms recorded by array element MK01 for the 32 basis events used to construct the rank-18 subspace detector. Although each of these detections exceeded the correlation threshold, the only commonality of these signals is the short P-pulse followed by low-amplitude coda. The TBW is apparently too low for these correlations to be significant. 34
Fig. 3.2.4 485 power detections (blue) superimposed on the 1222 aftershocks from the REB. These were used to create 4 subspace detectors with ranks ranging from 56 to 115. 34
Fig. 3.2.5 629 detections produced by the four high rank subspace detectors and associated to the REB events. 35
Fig. 3.2.6 All detections (shown on MK01) from the 10 days preceding the main shock. The traces in red are for detections with times that associate with events in LLNL database. The associated events are in the aftershock zone. 36

Fig.3.3.1 Map of Spitsbergen and the Storfjorden sequence. The event locations shown in green represent the NORSAR analyst-reviewed catalog between 21 February 2008 and 20 April 2012. Those in blue represent a subsequent relocation catalog (Pirli et al., 2013). The SPITS array is indicated by the triangle approximately 150 km NNW of the sequence. 38
Fig.3.3.2 Scatterplots of detections produced by the framework over a four year period indicate the results of the screening operation taken on detections to limit results to the Storfjorden sequence. The top plot depicts the distribution of events in velocity and time, the middle plot shows the distribution in azimuth and time, and the bottom plot the distribution of an FK quality measure (power in the FK spectrum normalized by power incident on the array). The red dots represent all events and the black dots represent those events passing the FK screen. 39
Fig.3.3.3 Distribution of framework detections in slowness space as a polar scatter plot. As in Figure 3.3.2, the red dots represent all framework detections and the black dots represent those detections that passed the FK screen. 40
Fig. 3.3.4 Histograms of the magnitude distribution for the NORSAR catalog (red) and the catalog constructed from framework detections (blue) show that correlators produced by the framework collectively have a significantly lower detection threshold than standard pipeline detectors. 41
Fig. 3.3.5 Cumulative frequency-magnitude distributions for the framework catalog (blue) and the NORSAR reviewed bulletin (red) indicate a completeness magnitude of 0.8 for the former and 2.1 for the latter. The framework effectively reduced the detection threshold by 1.3 magnitude units. 42
Fig. 3.3.6 Cluster lifespan plot (left panel) and number of events per cluster (right panel) demonstrate the episodic nature of activity in the source region. 43
Fig. 3.4.1 Known active mines in Western Kazakhstan and neighboring regions of Russia together with fully automatic single array location estimates of events in the region using only the ABKAR array. The locations of the sites are provided in Table 3.4.1. 45
Fig. 3.4.2 The twelve clusters of mining events containing 4 or more signals obtained in the time period 2008-001 to 2008-090 on the ABKAR array. A single spawning beam aimed at backazimuth 355 degrees and apparent velocity 8.0 km/s was used as a primary detector. 46
Fig. 4.1.1 Embedding of the ABKAR DPRK 2009 event recording in the ABKAR stream during the Kashmir aftershocks. The embedding ratio was 1:1, so that true relative amplitudes are preserved between observations of the aftershocks and the nuclear tests. One channel (ABK01) before and after embedding and filtered into the 1-3 Hz detection band is displayed. Clearly this embedded observation did not provide a challenging test of detection. However, the point of this exercise was one of classification, not sensitive detection, i.e. to determine whether the correct detector found the event. 49

Fig. 4.1.2 Embedding of the DPRK 2013 test in the ABKAR stream about 2/3 of the way through the 11-day Kashmir aftershock test sequence. A second small arrival was recorded about 300 seconds prior to the DPRK P phase in the embedded DPRK segment and was inadvertently introduced into the composite stream. The data have been filtered into the 1-3 Hz detection band in this plot. 50
Fig. 4.1.3 P beam recipe for the ABKAR array intended to represent a collection of detectors as might be found in a monitoring pipeline. The beam slownesses are spaced every 0.0625 sec/km, which is about half-width full maximum on the ABKAR array response at 2 Hz (see Fig. 4.1.4). The data were filtered into the 1-3 Hz band prior to beamforming. The red cross indicates the theoretical vector slowness for a P wave originating at the DPRK test site. The closest beam is at 60 degrees, 16.0 km/sec, but the beam that detected the two embedded events was the next closest at 90 degrees, 9.24 km/sec. This result is an indication of P wave refraction near the array. The black asterisk indicates the vector slowness of the spawning beam directed at the Kashmir aftershocks, which replaced the nearest recipe beam. 51
Fig. 4.1.4 Theoretical wavenumber response pattern for the ABKAR array, superimposed with the slowness samples of the beam recipe scaled to wavenumber at 2 Hz. 52
Fig. 4.1.5 Distribution of beam recipe detections by slowness in the baseline calculation. The intensity of color indicates the number of detections by each beam. 53
Fig. 4.1.6 The distribution of recipe beam detections with slowness when a correlation spawner directed at the Kashmir aftershocks was added to the framework configuration. The distribution changes very little apart from the substitution of the spawner-derived detections for the 124 detections from the recipe beam replaced by the spawner. For this scenario, the addition of a spawner does little to alter the performance of the detection framework. 53
Fig. 4.2.1 Map showing the location of the 2011-2013 Tohoku sequence and the three teleseismic arrays used to observe it: ASAR (Australia), ILAR (Alaska) and PDAR (Wyoming). The cloud of white symbols off the east coast of Japan shows the extent of the aftershock sequence. 55
Fig. 4.2.2 Five of the ten largest clusters produced by the subspace clustering algorithm. The white background symbols show the positions of all events in the NEIC catalog. Each cluster is given a unique color. Note that events are widely distributed in these clusters, probably a consequence of the multi-rank nature of the subspace similarity metric. 61
Fig. 4.2.3 The remaining five of the ten largest clusters. Cluster 62 is the compact group of events on the northeast coast of Honshu represented by black circles, and cluster 16 is the large scattered cluster of events denoted by the blue circles. 62

Fig. 4.2.4 Cluster 62, an example of a geographically compact event group. Three events are depicted, with channels AS01, IL01 and PD01 repeated in that order for each of the three events. Note the high similarity of the waveforms. This cluster is indicated by black symbols in Figure 4.2.3..... 63
Fig. 4.2.5 Waveforms of three events drawn from a geographically diffuse cluster, the group indicated by blue symbols in Fig. 4.2.3. These waveforms are much more variable, most noticeably at ASAR, the first station of each triple of waveforms. 64
Fig. 4.3.I Illustration of cross-correlating the signal on the KKAR array generated by the Kashmir main shock (starting time of 30 second long template: 2005-281:03.52.50.385) with the signal from a significant aftershock some 90 minutes later. Both of the signals have high SNR but clearly have very different forms and do not result in correlation coefficients significantly above the background level. It is, for example, not possible to measure the time-delay between signals using a straightforward cross-correlation, even using the array stack..... 66
Fig. 4.3.2 Left panels: narrow band f-k analysis when filtering with a center frequency of 2 Hz. In the right panels, the center frequency is 4 Hz. In the top panels we used data from the time segment which corresponds to P_n phase arrivals from the main event. In the bottom panels, a time segment associated with a significant aftershock was selected. 68
Fig. 4.3.3 Narrowband f-k grids for the first arrival of the aftershock generated after pre-steering by the empirical matched field steering vectors calculated from the main shock first arrival at 2 Hz (left) and 4 Hz (right). 69
Fig. 4.3.4 Narrowband f-k grids for the Sn arrival of the aftershock generated after pre-steering by the empirical matched field steering vectors calculated from the main shock Pn arrival at 2 Hz (left) and 4 Hz (right). 70
Fig. 4.3.5 The top panel pseudo-spectrogram shows the relative power matched field statistic Pc as a function of time and frequency. The waveforms centered on 1.5, 3, and 4.5 Hz display the KKAR Pn beam steered towards the source region, Butterworth band-pass filtered at the center frequencies indicated with bandwidth 1 Hz. The middle panel shows the STA/LTA ratio calculated for each frequency bin of the relative power pseudo-spectrogram displayed in the top panel. The length of the LTA segment is 5 seconds and it is separated from the 1 second long STA segment by a gap of 5 seconds. The bottom panel shows a black line for the relative power statistic Pr averaged over all frequencies of the top panel. Similarly, the orange line is frequency average of the STA/LTA ratio of the middle panel as a function of time. All times are UT on day 281 in 2005. In the top and middle panels, the vertical axis is the frequency in Hz. The vertical dashed lines correspond to arrival time estimates for events listed in the Reviewed Event Bulletin (REB) of the PTS. 71

Fig. 4.3.6 The top four panels are pseudo-spectrograms for the KKAR, BVAR, WRA, and ILAR arrays for a 1 hour sequence on day 281 in 2005. For each array, empirical steering vectors were calculated for the first arrival from the main shock (this is a regional Pn phase for KKAR and teleseismic P for the other three arrays). All of the traces are back propagated to the origin time for the main event. The bottom four panels show the corresponding frequency-dependent STA/LTA ratios calculated in a similar manner as in the previous Fig. 4.3.5. The vertical dashed lines correspond to arrival time estimates for events listed in the Reviewed Event Bulletin (REB) of the PTS. 73
Fig. 4.3.7 Scalar functions derived from the STA/LTA transformations of the matched field pseudo-spectrograms for seven seismic arrays as labelled. Each of the traces is shifted backwards in time by the traveltime to the station from the main event (from which the EMFP template was defined). Two sets of peaks in these functions have been associated and assigned to events labelled “Event 1” and “Event 2”. 74
Fig. 4.3.8 If we fix the location of Event 1 to the point indicated by the white circle and measure the times of the corresponding maxima in the matched field scalars, we can calculate contours of misfit for the location of Event 2 by summing the residuals between predicted traveltime difference and measured traveltime difference (essentially a double-difference relative location estimate). Even using only the seven global seismic arrays displayed in the previous figure, we are able to constrain the location of Event 2 using only the times of the local maxima of the matched field output. The current best available analyst location estimate for Event 2 falls approximately at the center of the region with the lowest traveltime difference residual. 75
Fig. 4.3.9 Panel displaying a matched field scalar trace at KKAR (in red) together with the actual array beam (above) for the October 23, 2011, M=7.1 Van earthquake in Eastern Turkey. For the panel to the right, two clear detections are observed in the matched field trace for which no signal onsets are visible on the beam for the same array. In the very top trace is a seismogram from a local station (at a distance of approximately 1 degree), corrected for the traveltime from the source region, which displays a signal corresponding to the time of the matched field P-phase detection at KKAR. In the lowermost two traces are cross-correlation traces with the signal from the main shock: although a modulation is observed, the peaks are not well-defined in time or significantly above the background level. 76
Fig. 4.4.1 f-k grids surrounding primary detector triggers on the ABKAR array for confirmed events from the Terensay mine (left) and the Chromtau mine (right). (Details of the locations of these mines relative to the ABKAR array are provided in Table 3.4.1 and Fig. 3.4.1) The red box denotes an approximate region of slowness space for which detections will be accepted for spawning pattern detectors for monitoring events at the Terensay mine. If f-k analysis indicates a direction not contained within this region of slowness space then the detection is discarded. 78

Fig. 4.5.1 The detection of spurious events and failure to detect a correlating event due to the contamination of the waveform template with a foreign signal within the signal template. A signal starting at a time 2005-283:10.48.35 (in trace b) correlates well with a later signal starting at a time 2005-283:14.31.15 (trace a). However, under the initial fixed-length template recipe, the master waveform also included a high amplitude signal beginning at a time 2005-283:10.51.13 (trace c). Correlating this template with the incoming waveforms results in the detection statistic (trace d) which completely fails to register the repeating event and produces a spurious detection later in the time-window. Selecting a 10 second shorter template (trace e) avoids the interfering signal and registers a significant correlation at the time of the repeating event signal. 80
Fig. 4.5.2 An exact replication of the correlation calculations displayed in Figure 4.5.1 except that all filtered waveforms have been scaled by a moving average of the absolute values prior to the correlation. In this case, the influence of the interfering signal is marginal and the repeating event signal is detected by both templates. Also, no spurious detection is made later in the data segment. 81
Fig. 4.5.3 The correlation traces over a long time interval for the 165 second long templates starting at time 2005-283:10.48.35 with and without a scaling by the moving average (in traces 1 and 2 from the bottom respectively). Only a single detection is made for the RMA trace whereas multiple detections are made when the correlation is performed on the unscaled traces..... 82
Fig. 4.5.4 Example 1 of setting a template according to the time center and time-variance. 83
Fig. 4.5.5 Example 2 of setting a template according to the time center and time-variance. 84

APPENDIX – List of Figures

Fig. A.1 Events from the bulletins of the International Seismological Center (www.isc.ac.uk, left) and the Reviewed Event Bulletin (REB) of the International Data Center (IDC) for the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) in Vienna (right) between October 8, 2005, and December 31, 2005. 90

Fig. A.2 Contours of 1-norm time-residuals for arrivals (both at regional and teleseismic distances) for a single event in the sequence (Event 1 in Figure 4.3.8) both with and without empirical time corrections solved for in this study. Using only teleseismic P-phases results in a location estimate close to that displayed in the right panel, but the addition of many uncorrected regional phases at stations with an unfavorable geometrical distribution relative to the source region pushes the location almost 20 km to the southwest and increases significantly the traveltime residual and hence the confidence in the event location estimate. Time corrections have been calculated for all stations but are far larger for the regional phases than for the teleseismic phases. When the time calibrations are imposed, the time-residual decreases greatly. The same time corrections are applied to all events in the sequence and are likely to be especially significant for smaller events for which we only have recordings at regional distances. 93

Fig. A.3 A comparison between events relocated in the current study (left) and locations from the REB (right). The colors indicate the delay in seconds between the first P-arrival at NIL to the south and KKAR to the north for both distributions of event location estimates. While this does not provide an independent quality check on the event location estimates (since both stations were used in the event relocation), there is clearly a far more consistent pattern for the relocated events..... 94

This page is intentionally left blank.

SUMMARY

A detection framework has been developed which augments a routine processing pipeline with mechanisms for spawning and administering so-called pattern detectors to identify and organize repeating waveforms discovered in multichannel seismic data streams. The autonomous classification of almost repeating seismic signals is necessary to mitigate the strain on analyst resources which occurs under extensive aftershock sequences following very large earthquakes, and when faced with the task of assigning vast numbers of anthropogenic seismic signals to known sources of repeating explosions. The framework has been tested and evaluated on a variety of different test cases from mining blasts in Central Asia, to moderate earthquake aftershock sequences recorded at regional distances, and to large aftershock sequences on the scale of those generated by, for example, the 2005 M_W 7.6 Kashmir earthquake. The framework performs exceptionally in identifying repeating mining blasts, with signals from large clusters of events belonging exclusively to known quarries being detected and grouped fully automatically. For the moderate aftershock sequence scenario, the framework provided an event catalog down to a significantly lower completeness magnitude than could be generated with the available analyst resources. The large earthquake scenarios are far more challenging given the exceptionally large source regions covered by the aftershocks. The spatial separations between the hypocenters of the largest events (i.e. those well recorded teleseismically) are simply too large to result in significant waveform semblance, in the frequency bands of interest, between the signals from subsequent events. In these cases, the framework often identifies clusters of events in pockets of seismicity within the extended source region: often at lower magnitudes than appeared in the existing catalogs. The formation of high-rank subspace detectors can sweep up large proportions of the seismicity although not necessarily in a way that reduces significantly the burden of analyst interpretation. We suggest the use of a subspace-measure of waveform similarity that may perform better than the classical correlation coefficient for forming and evaluating clusters in such cases. Submerging a signal from the DPRK nuclear test into data from the same array from an extensive aftershock sequence, we demonstrate that this important signal is not screened out by the pattern detector components of the framework. We propose empirical matched field processing (EMFP) on single array streams as a sensitive primary detector for generating triggers at, and only at, the times of arrivals from events in the region of interest. The matched field detector shows promise for generating output that may be processed incoherently over multiple arrays for the creation of robust event hypotheses over an extended source region.

ACKNOWLEDGEMENTS

We would like to thank staff at the Kazakhstan National Data Center for permitting the use of data from the KKAR and ABKAR arrays in this work and we thank Terri McDonald Hauk at Lawrence Livermore National Laboratory for preparing the data segments.

We would like to thank Bill Junek for his continued interest in the development of the detection framework, and for giving us permission to summarize his results regarding application of the detection framework to the Storfjorden aftershock sequence on Svalbard. These results, presented in section 3.3 of this report, partly reflect cooperative efforts conducted during his research visit to NORSAR in the spring of 2013.

PROJECT OBJECTIVES

The objective of this project is to determine whether processing pipelines can adapt in real time to classify aftershocks in major developing sequences. The intent is to catch aftershocks in the detection front-end of pipelines, possibly removing them from subsequent processing or subjecting them to efficient and streamlined processing in later stages of the pipelines. The goal is to ameliorate analyst overload during major aftershock sequences and to automate discovery and classification of repeating sources, to include mining explosions in addition to sequence aftershocks. The technical concept for achieving these objectives is an autonomous subsystem for creating (spawning), in real time, pattern detectors (array correlators, subspace detectors) responding to sequence aftershocks and other repeating events. Templates for the pattern detectors would be extracted from the stream by standard power detectors running in the pipeline making aftershock detections. These detectors would operate in addition to the standard recipe beam detectors and would be deployed to process data streams on-the-fly.

The engineering objective of the project is to construct a functioning model of the detection stage of a pipeline implementing conventional beam recipes, but extended to create and manage pattern detectors under a variety of spawning policies. This model system, referred to as the framework, will allow tests of several strategies for discovering repeating waveform patterns and organizing detected occurrences for efficient interpretation by analysts. The system is conceived to maintain an archive of metadata on detector creation, triggers and detections associated with all detectors, and all configuration parameters in a relational database. This archive is to allow

analysis of system evolution, since the system will be dynamic, and to allow processing to be resumed or recreated with edited versions of system state to test the influence of different spawning policies and parameter settings on overall performance. The system also is to implement an autonomous supervisory function that keeps track of detector performance, culling, updating or merging detectors to improve performance. This includes a function to periodically reprocess the data with a suite of mature pattern detectors to allow detection of early weak events with patterns from high-SNR events that occur later in the sequence.

The testing objective is to grade system performance by running the framework on aftershock sequences. The ultimate metric of performance is a measure of the consolidation of detections into efficiently-interpreted families. The principal test is an analysis of the 2005 Kashmir sequence with the four regional Kazakhstan arrays as a network.

Other detailed technical objectives include:

- Conducting tests to insure that the autonomous system does not screen events of interest from analyst evaluation, by incorrectly labeling explosions of interest as aftershocks or mining events.
- Conducting tests of the potential for extending pattern detection to networks of arrays to obviate problems in association stages with building incorrect events.

1 INTRODUCTION

Despite years of research and development on seismic network operations and decades of practical experience with processing pipelines, aftershock sequences associated with large earthquakes (magnitude 6 and above) still overwhelm analytical resources. Very large events such as the 2004 Sumatra-Andaman Islands earthquake, the 2008 Sichuan earthquake and the 2011 Tohoku earthquake flood seismic networks with thousands of aftershocks in the weeks that follow the main shock. Aftershocks may be distributed over large geographic regions.

One strategy proposed to cope with such challenges exploits the fact that aftershock sequences often produce events with highly similar waveforms. In such circumstances, it may be possible to group events on the basis of waveform similarity and structure analytical work around such groups to reduce work load. Several approaches have been suggested. The most ambitious would limit analyst interpretation to a small sample of events that represent the range of waveforms produced by a major sequence or swarm; perhaps only one or a few events in each group. For this approach to be trusted, events must be grouped on the basis of highly reliable measures of waveform similarity. Previous experience at near-regional distances and with smaller aftershock sequences suggested that this strategy might be viable.

A more conservative approach would map automatically an analyst's interpretation (picks, source location, etc.) from a master event to events subsequently detected by a correlator generated from the master event (Junek *et al.*, 2013). All events would be reviewed in this approach, but the analyst would be given a leg-up with a correlation-determined initial interpretation for newly-detected events. This approach also could provide hints to (or place constraints upon) the automated association process building events from detections across a network. Again, these hints would be provided from an analyst reviewed association of master event arrivals across the network.

To be useful, a system built to realize either of these strategies must discover and associate repeating waveform occurrences in real time as the sequence unfolds.

Correlation, subspace and matched field detectors (here collectively referred to as *pattern detectors*) provide efficient means to detect and associate occurrences of repeating waveforms, and may form the basis of a near-real-time system to group and screen aftershocks. These detectors are constructed from waveform templates derived from type (master) events. To develop and deploy them in most implementations is

labor-intensive, if done manually. The aim of this project has been to examine possible extensions to existing pipeline systems that automate creation of pattern detectors to detect and organize waveforms from a developing aftershock sequence.

Another conceivable use of pattern detectors is to obviate association problems that often cause pipeline backlogs during aftershock sequences. The association stages of pipelines are responsible for building events; often events are built incorrectly and must be manually broken apart and then rebuilt by an analyst. By extending waveform correlation calculations to multiple stations in a network, it may be possible implicitly to associate seismic phases as part of the pattern matching operation.

To examine the feasibility of extending existing pipelines with autonomous components to create and deploy pattern detectors in real time, we constructed a test framework that models the detection front-end of pipelines. This framework considerably improves upon an earlier test system described in Harris and Dodge (2011), principally by supporting coherent operations across networks of arrays or single stations and implementing many different kinds of detectors in a unified manner. The system creates (spawns) correlation and subspace detectors autonomously based on events detected by more traditional power detectors: STA/LTA detectors operating on array beams or three-component station traces. Hundreds to thousands of spawned correlators can operate in near-real time in the system due to careful attention to efficient signal processing and concurrent implementation. The system also archives, in an Oracle database, the full history of every detector, including circumstances of creation, design parameters, triggers and associated detections. The archive allows system evolution to be parsed and examined after the fact and to be reproduced. It also allows new runs to be initiated from the edited and modified final state of a previous run. In this latter capacity, the framework has been supplemented with auxiliary interactive software that provides a convenient interface for examining detection results and editing state (deleting poor-performing detectors, for instance).

We experimented successfully with a variety of policies and techniques to prevent runaway proliferation of automatically-created detectors. Broadly, two approaches appear fruitful. The first is to replace the traditional power detectors used to spawn correlators with detectors less likely to trigger on unintended events or noise bursts. We studied empirical matched field processing (EMFP) detectors for this purpose and found them to be significantly better at selecting events from a target aftershock sequence for use in spawning correlators than traditional beam power detectors. The second approach is to provide a battery of screens on the waveforms that trigger power detectors to ensure that those waveforms have characteristics consistent with

anticipated aftershocks. We found that wideband FK screens are highly effective in this capacity, and that simple checks on waveform duration and kurtosis also help to limit undesired detector creation.

The framework has been tested in this project principally, as proposed, on the 2005 Kashmir earthquake and aftershock sequence. However, the framework has been found by several institutions to be useful for examining a variety of issues surrounding aftershock sequences. It has been employed for research in other projects and the results of those projects contribute to the conclusions that we report here. Specifically, it has been used to study the 2008 Storfjorden sequence at Spitsbergen and the 2012 Sumatra earthquakes.

The principal conclusion of this project (and the others that have used the framework) is that successful application of correlation classification is strongly a function of the size of the aftershock sequence and the distance of observation. Paradoxically, the very largest events, which produce the most troublesome aftershock sequences at observation ranges greater than 10 degrees, do not produce enough well-recorded aftershocks to support productive grouping. This paradox appears to occur because the geographic area of the aftershock sequence grows faster with increasing magnitude than the number of aftershocks producing good observations. At teleseismic ranges, the events must have magnitudes near 5 and above to produce usable waveform templates. As the area enlarges, these events become farther apart on average. As is well known, waveform correlation declines rapidly with increasing source separation. The net result is that large aftershock sequences observed at great range have fewer "correlative twin" events than smaller aftershock sequences observed at closer range. Consequently, augmenting pipelines with autonomous correlator-spawning subsystems is likely to be worthwhile only for stations at near-regional distances.

Two studies of sequences associated with smaller main events indicate that a spawning framework can be effective at reducing work load or developing a more complete catalog of seismicity at low cost. A study of the 2003 San Simeon earthquake in a prior BAA project showed possible reductions of 40-70% in analyst work load in screening aftershocks (depending on how work load is interpreted) with the system that preceded the test framework of this project. An independent study of the 2008 M_w 6.2 Storfjorden earthquake and subsequent sequence (Junek *et al*., 2014, in preparation) using the current framework vastly increased the number of events that could be used to interpret this sequence at modest cost. The mapping of event clusters constructed by the framework exposed previously unknown spatio-temporal structure in the sequence and changed the interpretation of the underlying tectonic mechanism.

Our study also helped lay to rest one concern regarding autonomous correlator spawning: that the system might create detectors that incorrectly trigger on events of interest not associated with the aftershock sequence and cause them to be classified as aftershocks and missed. To test this possibility, we embedded the 2009 and 2013 North Korean nuclear test waveforms recorded at ABKAR in ABKAR data containing the first 10 days of the Kashmir aftershock sequence. We configured the framework to simulate a beam recipe for ABKAR with a dedicated spawning detector targeting the aftershock sequence. None of the spawned correlation detectors triggered on the nuclear test events, but an appropriate recipe beam detector did trigger on these events.

Finally, there is another motivation for building a system to discover repeating sources: detection of groups of explosions. Among other purposes, these find use in calibrating and testing discriminants. In some regions of the world, mining explosions dominate seismic detections. In such cases, automatically-created pattern detectors may perform the function of discriminating mine shots. The framework was successful in discovering groups of mining explosions in Kazakhstan using data from the ABKAR array.

2 THE DETECTION FRAMEWORK

2.1 Framework Architecture

Fig. 2.1.1 below, drawn from the project proposal, is a block diagram of the detection processor that we proposed to build. It is sufficiently close to the system that was realized to use as a basis for describing the framework. Our intention was to model the detection front-end of existing pipelines, but augmented with a facility for autonomously creating and running "pattern detectors" on a stream of network data. The pattern detectors, principally array correlation and subspace detectors, were intended to "sweep up" events from an aftershock sequence in a process of simultaneous detection and classification that would simplify the subsequent association process and reduce the interpretation load for analysts.

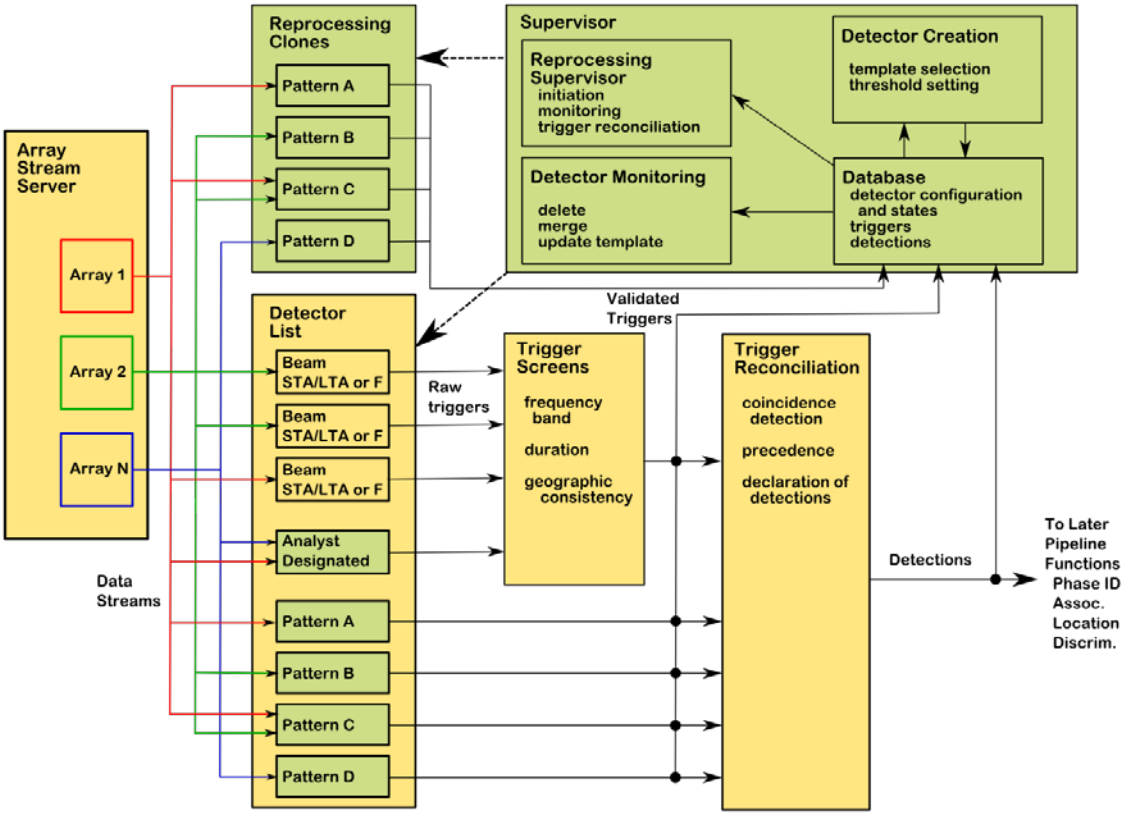

Fig. 2.1.1 The high-level block diagram of the detection framework as proposed and actually built. The boxes in yellow indicate functions approximately shared with a conventional pipeline detection processor. Those in green represent entirely new functions. Note that some detectors may process data from more than one array, possibly coherently.

The framework consists of six major functional components, which we describe in turn.

2.1.1 Array Stream Server

Existing pipelines have facilities to acquire streams of data from arrays and simple stations over telecommunication links and present them in an organized and standard form to detection pipelines. The framework emulates this process with an Array Stream Server that acquires data from flat files through a database interface. The signal processing routines used in the Stream Server maintain state between consecutive, contiguous blocks of stream data, in a manner that is consistent with continuous acquisition (meaning that it could be adapted for real-time operation). The principal difference between the framework and existing pipeline data acquisition lies in the fact that the Stream Server performs a global synchronization of the disparate streams from the individual arrays and stations that comprise the network. This means that the data on all channels are resampled to a common sampling rate and, beyond that, to common sampling time instants.

The purpose of global synchronization is to support the possibility of coherent processing across the entire network of stations, i.e. to make it possible to turn a network into a giant array.

The Array Stream Server can assemble multichannel streams from an arbitrary set of channels in a continuous-data CSS3.0 archive, implemented with an Oracle database. For example, a multichannel stream could consist of all vertical channels from three teleseismic arrays, e.g. ASAR, ILAR and PDAR. The Server provides fully synchronous, multichannel data in consecutive, contiguous blocks to the collection of detectors.

2.1.2 Detector list

The framework maintains a configurable list of detectors of many different types. The detectors have been freshly implemented for consistency, sharing a common code base in a hierarchical object-oriented fashion.

An unusual feature of our implementation is that detectors can be added or removed as the streams are being processed. As described above, the data streams are processed in consecutive, contiguous blocks. Changes to the set of detectors can be made at the conclusion of processing each block and before processing begins on the next block. At this point, new detectors can be created and added, and old detectors, perhaps not performing well, can be removed. All detectors, new or old, then immediately operate on the next block. One of the concepts we have tested with this arrangement in the

past, is to create an ecosystem of competing detectors, along with rules for selecting winners among the detectors. This flexible architecture allows research on options ranging from a fairly conventional collection of array beam recipe detectors, augmented with correlators, to an adaptive, self-optimizing pipeline.

The framework currently implements the following types of detectors:

- single or multichannel STA/LTA detectors, with incoherent stacking among channel outputs when multichannel
- array beam with STA/LTA detector on the beam
- array subspace detector (two types), with rank-1 detectors implementing a standard array correlator
- array matched field detectors (incoherent and coherent, with many options: prewhitening, multi-rank among them)

The notion of coherent array processing in these detectors, particularly for the empirical detectors derived from observed waveforms (subspace, matched-field) can be extended to arbitrary networks.

Among the detectors, the array beam power (STA/LTA) and matched field detectors can be designated as spawning detectors. This designation means that whenever one of these processors detects an event and there is no coincident detection from a higher-rank detector, the detected multichannel waveform, if well-recorded, is used to create a new array correlation detector. This correlator is immediately added to the detection list and begins operation with the next available block of data

One of the principal configurations that the framework supports, and which we have tested in the course of this project, is an array beam targeting an aftershock sequence and designated as a spawner. The intent of this configuration is to create a suite of correlators that detect and classify aftershocks, allowing some form of efficient aftershock processing. This configuration, in addition, has been augmented with a large ensemble of standard beam recipe detectors to examine whether a spawning detector and its progeny, intended to sweep up aftershocks, unduly interfere with the operation of the recipe detectors (they don't in the one scenario tested).

The framework has a two-step method for declaring detections. At the conclusion of processing for each block, the statistics produced by the detectors are scanned for two conditions, depending on the type of detector. Triggers are declared when detection statistics exceed a threshold, and the timing of the trigger is determined in one of two

ways: either by the point at which the statistic crosses the threshold (the method for STA/LTA detector types) or the point at which the statistic achieves its maximum in some neighborhood of the threshold crossing (the method for correlation and subspace detectors). Triggers among all detectors are compared in order to determine which should be promoted as detections (in the event that multiple detectors trigger simultaneously; see trigger reconciliation below).

Subspace (correlation) detectors also can be updated (optionally). When a subspace detector makes a detection, the newly-detected waveform can be used to update the collection of left singular vectors that define the subspace template. The purpose of this facility is to allow subspace detectors to track an evolving source. It has also been used to create high-rank subspace detectors in an attempt to increase the geographic footprint of a source region characterized initially by a correlator.

2.1.3 Screens for triggers

There was a concern that triggers on noise bursts, dropouts and signals from sources very local to the stations might produce runaway detectors that would flood the system with thousands of useless detections. As one approach to avoiding this situation, the framework has a facility for applying an arbitrary number of tests to screen undesirable triggers, particularly those from spawning detectors, to prevent their use to create pattern-matching templates.

The available screens include tests for minimum duration and bandwidth, and on direction and velocity (FK screens). The FK screens are highly effective at removing sidelobe triggers on the small arrays of the 9-element CTBT type that generally have high sidelobes.

We have found that FK screens are useful even for correlation detections when the thresholds for these detectors are set aggressively low.

Several other parameters associated with the trigger segments, like SNR, amplitude, time centroid, kurtosis, time variance, frequency variance, and number of glitches are calculated and stored in the database. These parameters are intended for further refinement of the classification of the triggers.

In addition to trigger screens, we examined the possibility of tracking detectors for runaway behavior and removing them from the list when found.

2.1.4 Trigger reconciliations

Often it is the case that multiple detectors trigger on well-recorded events. In considering the design of the framework, we decided on a principle that detections would be associated with ("owned by") a single detector. Consequently, when several detectors trigger nearly simultaneously, it is necessary to choose among the detectors which will own the event. The framework implements a set of rules for reconciling triggers and determining which trigger will be "promoted" to detection status and associated with its originating detector.

The rules implement detector precedence: simple STA/LTA detectors have the lowest precedence, array beam detectors have the next lowest precedence, matched field detectors have the next higher precedence and correlation/subspace detectors have highest precedence. Triggers from two or more detectors of the same precedence are reconciled in favor of the detector with maximum detection statistic.

The determination of simultaneity among triggers requires careful calibration of the timing of detection statistic threshold crossings or peak excursions. STA/LTA threshold crossings may occur at significantly different times than the peak of a correlation detector (it is comparing apples and oranges). Consequently we use empirically-determined rules for correcting the raw trigger times of all detector types to the threshold crossing of an array beam STA/LTA detection statistic. Simultaneity then is determined by comparing the corrected trigger times within a user-defined tolerance.

2.1.5 The Supervisor

The framework is highly evolved to collect diagnostics on the performance of the detectors that it operates and to make decisions and take action on that information. The archival infrastructure is implemented in an Oracle database with a custom schema. The framework maintains, in the database, an archive of all detectors including their configuration parameters, state information and a complete history of triggers and detections. For the empirical detectors (spawned subspace), the framework also maintains information on the parent spawning detector, design event and modification history (in case template update has been enabled). The trigger and detection data include originating detector, trigger times and detection statistic values.

The complete archive of detector construction and processing history allows post-mortem analysis of system behavior and recreation of processing runs with slight modifications to assess tuning changes. Given that system behavior is dynamic (and, by definition, unpredictable in detail), a post-processing examination process allows

pathological behaviors to be detected, characterized, diagnosed and their evolution dissected. New policies then can be designed and implemented to prevent their recurrence.

Initially we thought to have these archival and several other supervisory functions (indicated in Fig. 2.1.1) implemented in a separate Supervisor module. In practice, they have been distributed throughout the hierarchy of classes that comprise the framework.

The principal purpose of the framework is to implement and test a detector spawning function. This function examines detections made by designated spawning detectors and, if the detection waveforms pass a series of quality checks, creates and implements new correlation detectors. Typically, spawning is from one or more array-beam detectors targeting the P waves of an aftershock sequence. The framework also implements empirical matched field detectors, which may be designated spawners. We have conducted a significant amount of research on the use of matched field detectors that indicates they are more reliable than array beam power detectors for targeting aftershocks. When a new correlator is created from a spawning detector of either type, a variety of rules and waveform measurements (e.g. of template waveform duration) are invoked to determine the configuration parameters for the new detector.

The framework monitors detectors for pathological behaviors: inactivity and runaway behavior. Dead detectors make no detections once created and may eventually be pruned from the list. Runaway detectors exhibit the opposite behavior, making an inordinately large number of detections (the number is a user-defined parameter). This situation can occur, for example, if, despite the quality checks on triggering waveforms, a spike or other very short signal has been used to construct a correlator. When this behavior is detected, the offending detector is removed from the detector list.

2.1.6 Reprocessing clones

A separate clone processor can reprocess the data stream in parallel with the principal detection processing function. During reprocessing, the entire suite of detectors created since the last reprocessing time is cloned and re-initialized at the beginning of the sequence. This copy of the framework is run with those detectors and any triggers are reconciled among those detectors. When the reprocessing framework catches up to the point in time at which it was spawned, it is cleared and awaits the next reprocessing task. Note that while reprocessing is occurring, the detectors that were cloned for reprocessing continue to make progress in the main thread. There are no spawning detectors present in the reprocessing framework instance, so no new detectors can be created. Currently, the behavior is undefined if subspace updating is turned on.

There are two purposes for this facility: the first, discovered during early work on detector spawning systems, is that good template events may occur late in an aftershock sequence. Early events producing the same waveform pattern may be missed by array power detectors, due to low signal levels or occurrence in the coda of a much larger event. Reprocessing the streams with later-defined detectors retrieves these events for inclusion in a more complete catalog.

2.2 Research Environment – Supplementary Software

The framework has evolved into a flexible testbed for trying a variety of different policies for creating correlation and subspace detectors, and has become a useful tool at 3 institutions for conducting research on aftershock sequences and other repeating sources. One consequence of this wider usage (beyond this initial project) is that the framework now requires a large number of parameters for its configuration. Consequently, several supplementary codes have been developed to simplify the complex tasks of preparing the framework configuration for a particular test and interpreting the results. Fig. 2.2.1 is a schematic of the workflow for research using the framework. A family of four programs has evolved to support that workflow.

A configuration creation program, *ConfigCreator*, simplifies the creation of parameter text files which collectively contain the specification for a framework execution ("run") and subdirectories for storing detection statistics and other waveform data (Fig. 2.2.2). The configuration files determine the number of detectors, their types, all parameters, the details of the data streams they operate on (which arrays, channels) and preprocessing parameters (filter cutoffs, decimation rates).

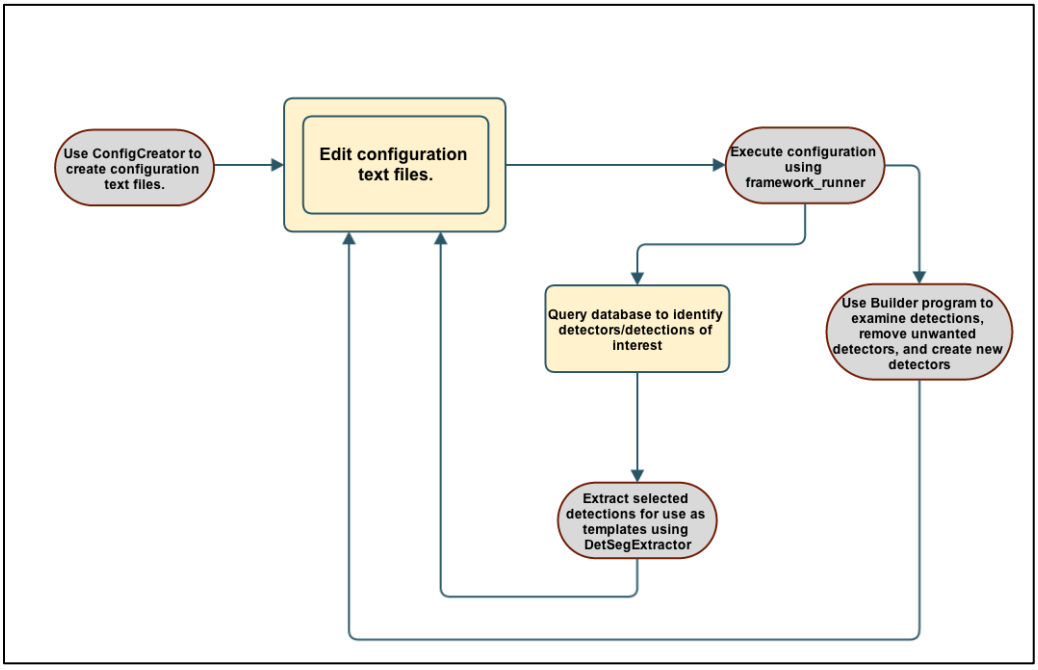

Fig. 2.2.1 Simplified schematic of the research workflow surrounding execution of the framework. As the framework has evolved to support a number of research objectives at 3 institutions it has become more flexible and complex, with a large number of configuration parameters. Its operation can store the triggers and detections associated with thousands of detectors. Four tools have been developed to manage the workflow.

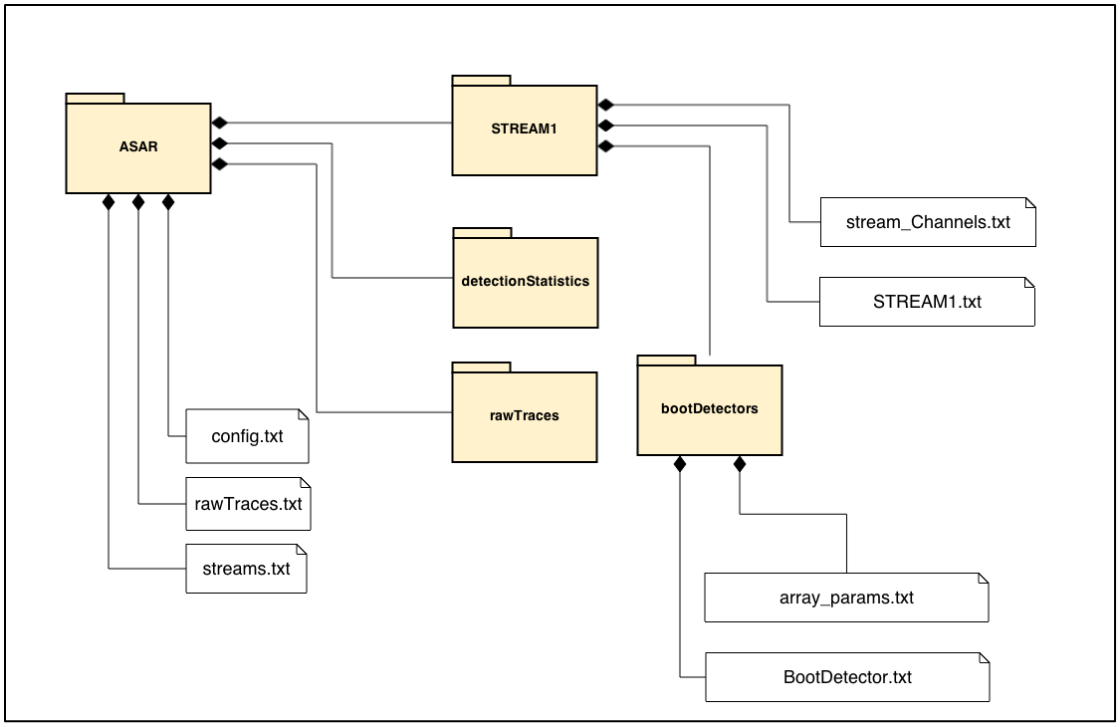

Fig. 2.2.2 Subdirectories and parameter text files created by ConfigCreator to support the execution of the framework. The framework actually is run using the command-line script framework_runner.

The framework is run with the command-line script *framework_runner* which is invoked with the config.txt file that contains references to the hierarchy of parameter files required to specify the full execution environment. Among the parameters is one that causes the framework to display, as it is running, a graphical user interface (Fig. 2.2.3) showing the waveforms of each detection as the detection is made. This feature often provides insight into unforeseen dynamical behavior of the framework as new detectors are created.

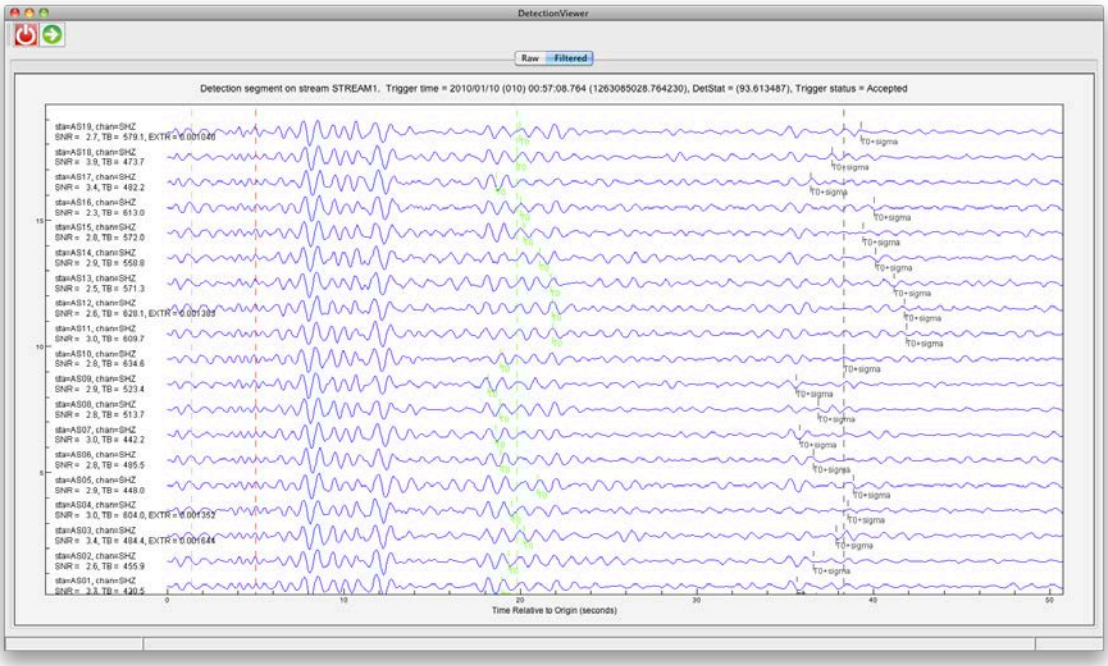

Fig. 2.2.3 Screenshot of the diagnostic GUI optionally displayed as the framework is running.

The Builder program (along with queries to the database which the framework uses to maintain state and store results) simplifies interpretation of results and provides the means to manage complex configurations with hundreds of detectors. Often the framework is run multiple times to bootstrap a large family of detectors to characterize a particular aftershock sequence. A first run of the framework might for this purpose initially implement just a single array power detector to spawn correlators. A second run then may be initiated with the correlators created in the first run, after the *Builder* has been used to prune away poor performers or detectors exhibiting some pathology (runaway behavior, templates with two superimposed events, etc.).

The Builder has a graphical user interface used to review the waveforms of detections associated with each detector (Fig. 2.2.4). The GUI is organized by detector ID – by clicking on the ID (at left in the figure), a view of the waveforms corresponding to the detections associated with the detector is shown in a separate panel (at right in the figure).

Finally, the environment has a command-line tool (*DetSegExtractor*) for extracting waveforms to SAC files for detections associated with a particular detector. This facility allows further examination of detected signals with a wider variety of analytical tools (SAC, EP, Matlab).

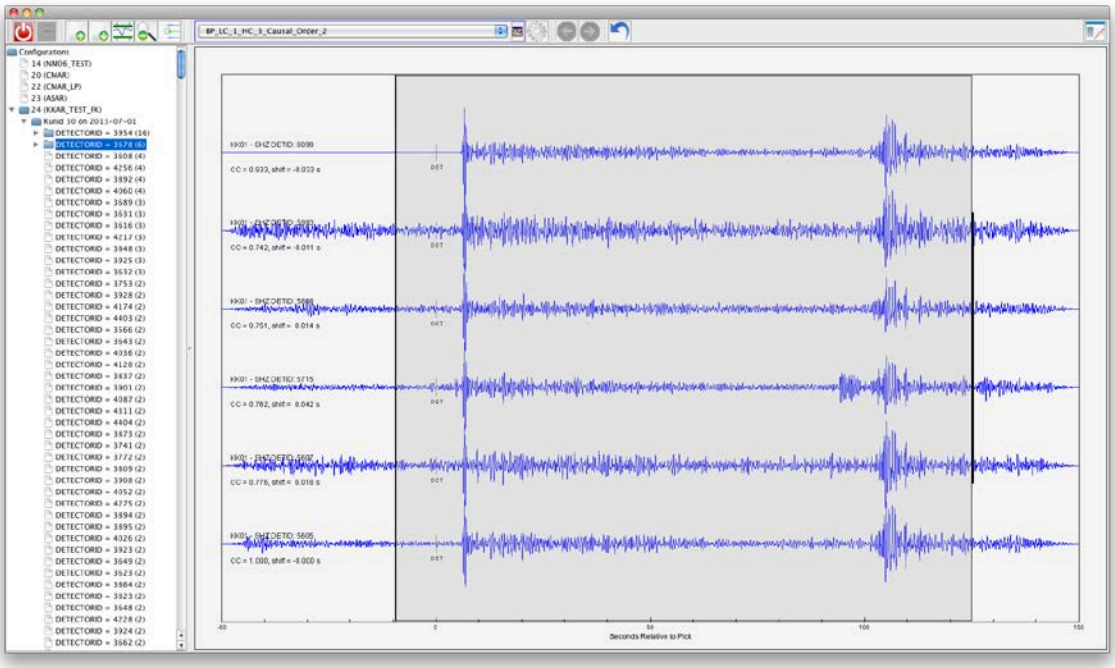

Fig. 2.2.4 Screenshot of the Builder GUI used to review waveforms of events detected by a particular detector. The detector is selected from the list displayed in the left panel and the event waveforms appear in the right panel.

3 EVALUATION OF FRAMEWORK PERFORMANCE

The work plan for this project included evaluation of the detection framework and the concepts it embodies on a representative aftershock sequence, associated with the 2005 Kashmir earthquake. Since the inception of the project, the framework has, in addition, found application elsewhere, supporting other projects funded separately from this one. Consequently, the automated spawning techniques implemented in the framework have been tested in several contexts in addition to the Kashmir sequence. The collective results of these investigations provide insight on the general topic of correlation detection as an aid to reducing analyst work load.

The general picture that has emerged is that the strategy of creating ensembles of correlation detectors to sweep up aftershocks is successful primarily with sequences associated with smaller main shocks observed at relatively close (near regional) range. The framework approach also appears successful for automating the detection and characterization of mining explosions, again at near-regional distance. However, for sequences associated with great earthquakes, observed at distances greater than ten degrees, well-recorded aftershocks are not sufficiently dense in the source region to provide a comprehensive set of waveform templates. While other strategies involving more complex subspace detectors may yet provide a means for "covering" a large source region with suitable templates, the first order approach of spawning simple correlators from array power detections appears inadequate to address the classification problem for great events.

In this section, we report results of four studies that used the detection framework, two undertaken by this project and two conducted by separately funded projects. The first (Kashmir sequence) and last (mining explosions in Western Kazakhstan) were part of this project. We summarize results from a project carried out by LLNL under U.S. State Department funding investigating the 2012 Sumatra earthquake sequence and another project carried out under U.S. Air Force funding on the 2008 Storfjorden sequence. Collectively, these studies span a range of main event magnitudes and observation ranges.

3.1 The 2005 Kashmir sequence

The aftershock sequence associated with the M_w 7.6 earthquake of 8 October 2005 recorded at the Kazakhstan arrays, was proposed as the primary test case for evaluation of the detection framework developed within this project. We here report on the results from processing of aftershock data from the 9-element KKAR array, located about 9° north –northwest of the earthquake source area (see Fig. 3.1.1).

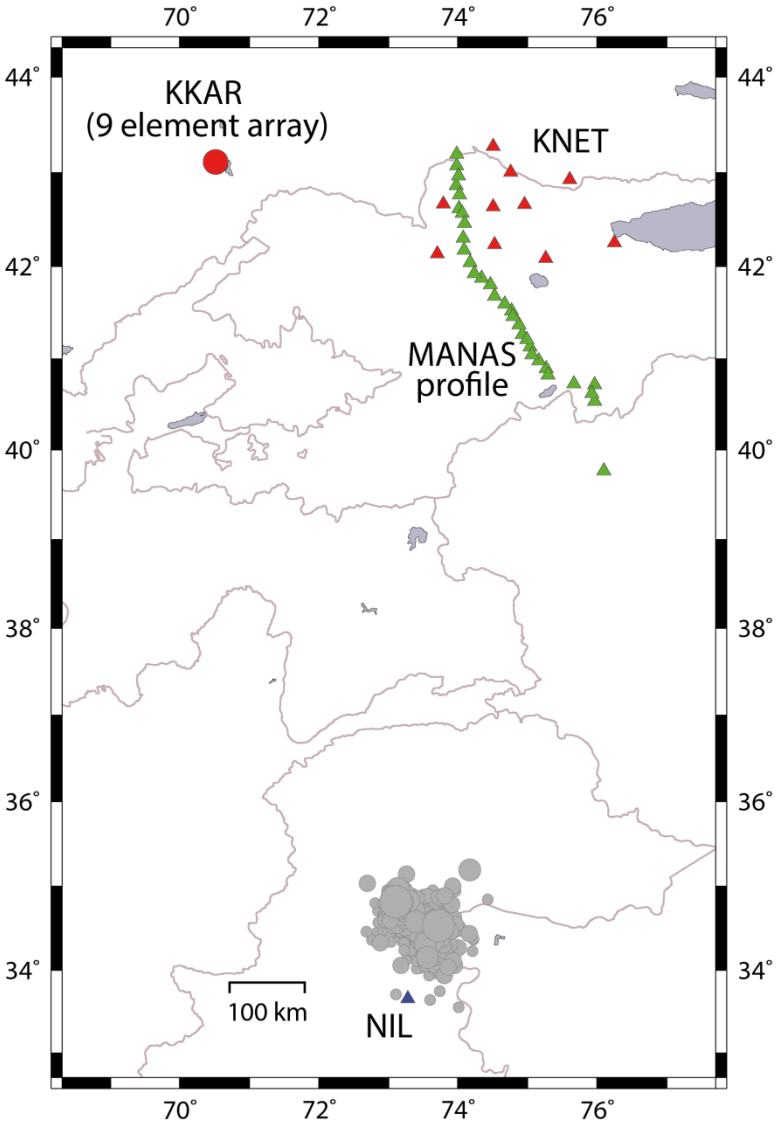

Fig. 3.1.1 Map showing the IDC Reviewed Event Bulletin (REB) locations of the 2005 Kashmir earthquake aftershock sequence (grey symbols) together with KKAR array (red filled circle) and other stations in the region recording the events (triangles).

To obtain a high-quality reference event database for evaluation, available regional and teleseismic data were reanalyzed, and rather extensive relocation efforts were made to improve the quality of the events locations. This effort is described in the Appendix to this report, and resulted in a reference set of 786 events during the 10-day period. Figure 3.1.2 shows recordings of three different aftershocks at one of the vertical-component sensors of the KKAR array. Notice the difference in signatures of the waveforms, in particular regarding the amplitude ratio between the P- and S-phases.

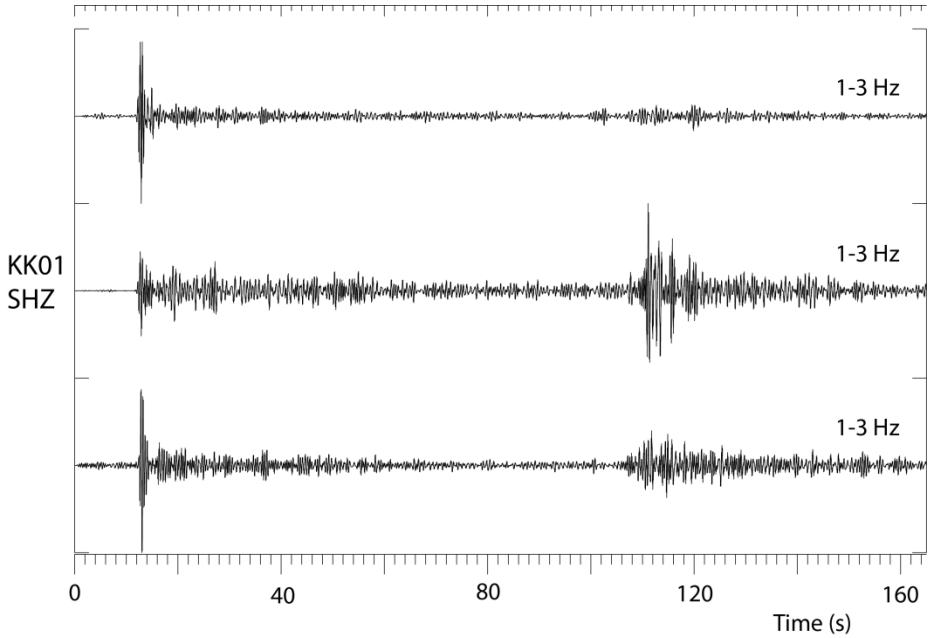

Fig. 3.1.2 Recordings of three different Kashmir aftershocks at the KK01 vertical component sensor of the KKAR array. The data are bandpass filtered between 1 and 3 Hz.

3.1.1 Running single event correlators at KKAR

In order to evaluate the potential of the detection framework for creating clusters of events from Kashmir aftershock sequence, we processed KKAR data for a 10-day interval starting at the day of the main event (8 October 2005). The initial triggers were made on a filtered (1-3 Hz) array beam steered with the approximate back-azimuth (160°) and apparent velocity (9 km/s) of the P-phase of the main event. A linear SNR of 5 was required for detection, and an FK screen was applied to validate the primary detections (triggers). SNR was defined as the STA/LTA ratio of the beam amplitudes. The back-azimuth and apparent velocity tolerances of the FK screen were ± 5° and ± 1km/s, respectively.

For each new initial P-phase trigger, a rank-1 subspace detector (correlator) was launched for processing by the framework (see section 2.1 for details). To include both the P- and S-phases, the default template length was set to 150 seconds. However, to accommodate situations with *e.g.*, overlapping events, the template lengths could be reduced down to 50 seconds. This is done automatically within the framework (see Section 4.5). A bandpass of 1-3 Hz was applied to the array data, and a linear threshold of 0.55 was required for detection of a new event by the correlators.

Somewhat surprisingly, only 4 clusters with more than 3 events were found by the detection framework during the 10 day period of the aftershock sequence. The waveforms of the largest cluster, and the corresponding event locations are shown in Fig. 3.1.3. We observe that the characteristic waveform of this cluster consists of an impulsive P-phase, relatively weak P-coda, and a weak S-phase. The strong energy of the P-phase heavily weights the calculated cross-correlation coefficient, effectively reducing the time-bandwidth product of the waveform template, and the relatively wide spread of the events over the source area is interpreted as being the results of high similarity between the P-arrivals only.

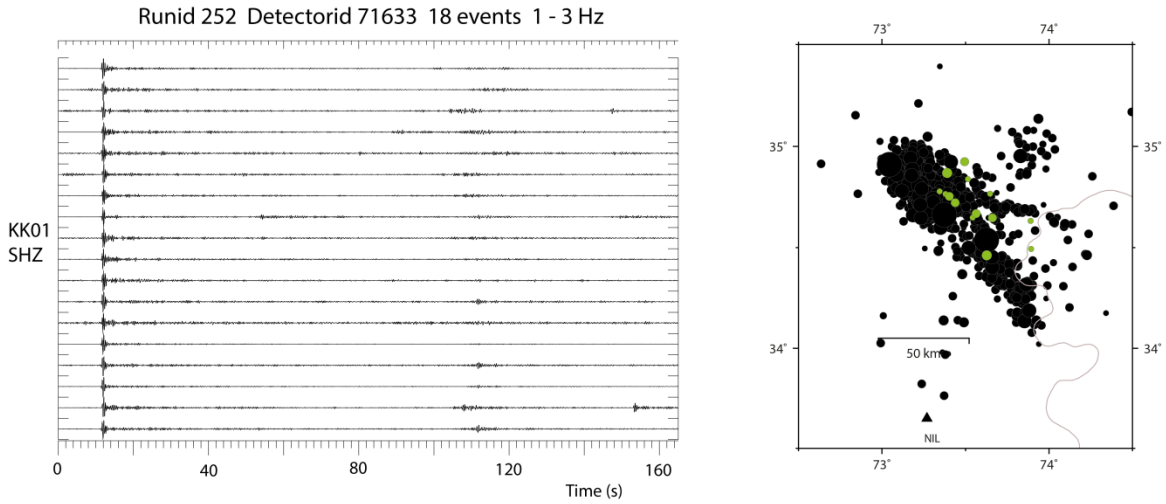

Fig. 3.1.3 Left: Waveforms at the KK01_SHZ sensor of the 18 events of the largest event cluster. Right: Locations of the 18 events (green filled circles) shown within the locations of the remaining events in the reference database (black filled circles).

The second largest cluster consisted of 7 events, and the waveforms and locations are shown in Fig. 3.1.4. The characteristic waveform of this cluster consists of both strong P- and S-phase, which are both weighted into the calculated cross-correlation coefficients. In this case, only events located within a small geographical exceeded the correlation threshold of 0.55.

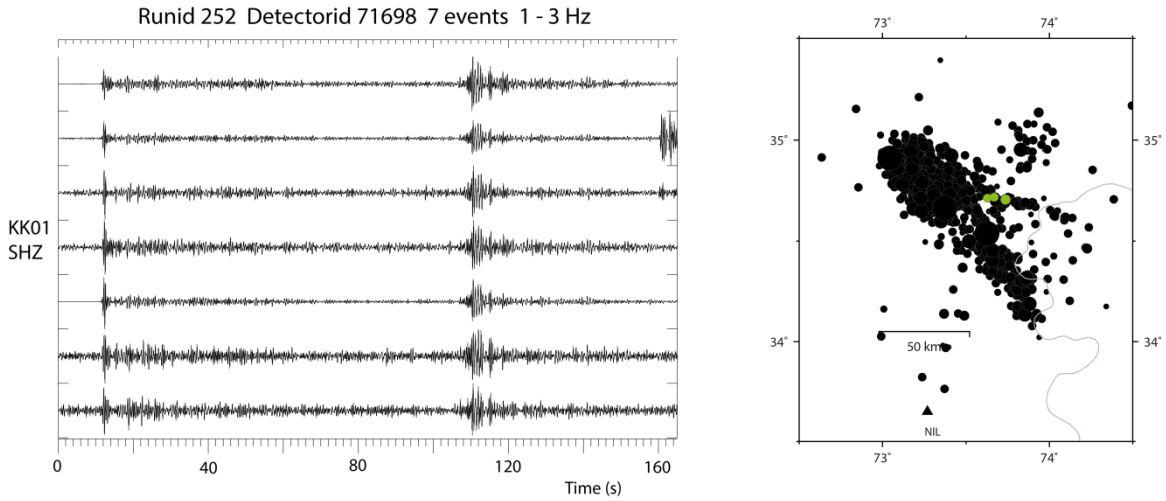

Fig. 3.1.4 Left: Waveforms at the KK01_SHZ sensor of the 7 events of the second largest event cluster. Right: Locations of the 7 events (green filled circles) shown within the locations of the remaining events in the reference database (black filled circles).

To shed more light on the potential of clustering of the aftershock sequence, we calculated for KKAR the mutual array-based cross-correlations among all 786 events in the reference database, and ran a clustering algorithm. The results are shown in Fig. 3.1.5 in terms of the mutual correlations ordered after cluster analysis. We observe a relatively small number of well-defined event clusters, which is in accordance with the previously described processing results from the framework. This suggests that single event correlators may be ineffective for detecting and classifying larger numbers of events from aftershock sequences associated with major earthquakes, for which the source region usually spans a large areal extent.

This observation is also in accordance with the findings of Slinkard *et al.* (2013), where they analyzed the cross-correlations of Kashmir aftershocks observed at the local station NIL in Pakistan and the regional station AAK in Kyrgyzstan and found very few repeaters.

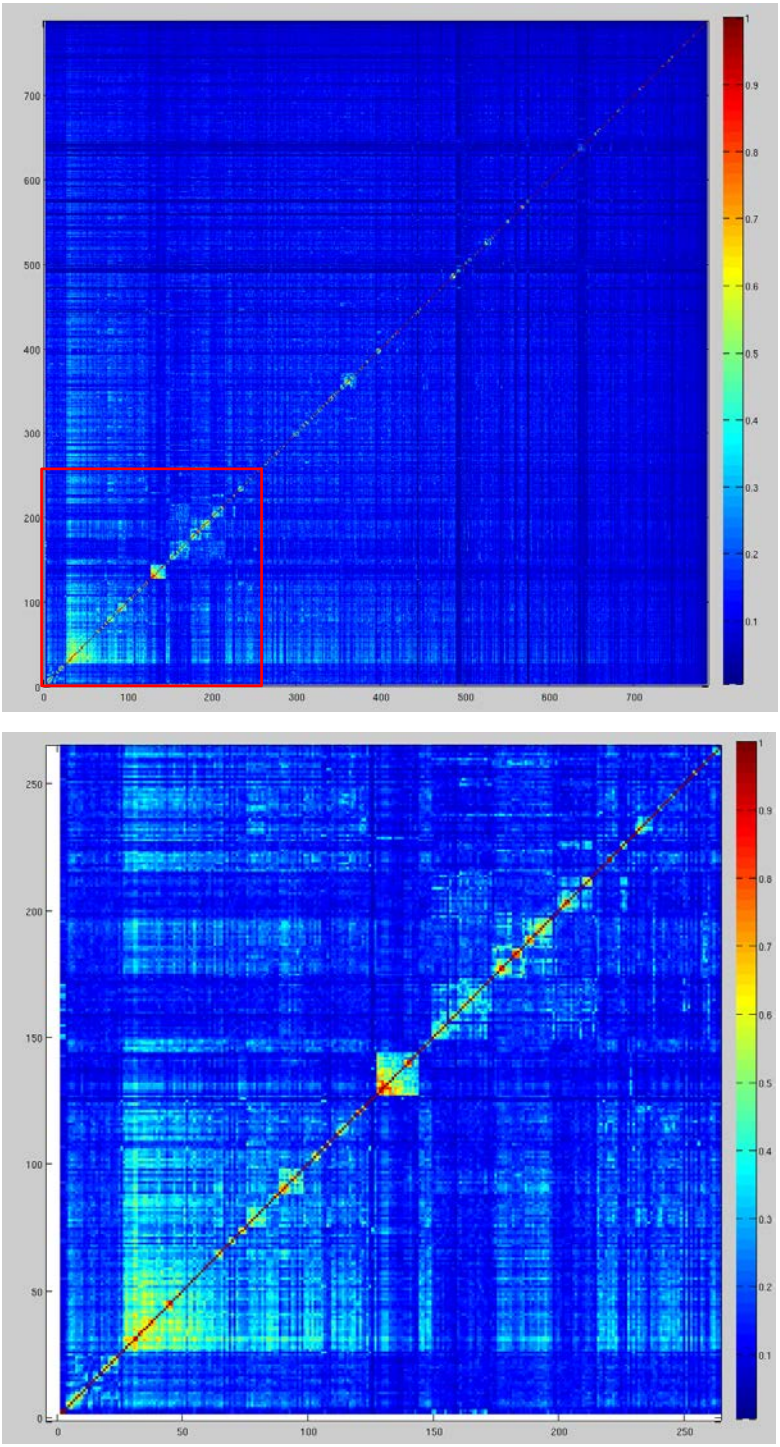

Fig. 3.1.5 The upper panel displays the mutual KKAR cross-correlations among the 786 events in the reference database ordered after cluster analysis. The scale ranges linearly between 0 and 1. The lower panel shows a zoom-in on the first 260 events within the red box of the upper panel.

3.1.2 High-rank subspace detectors

The limited ability of rank-1 subspace detectors (correlators) to form clusters which cover such a sequence to a significant extent motivates the building of higher rank subspace detectors. Constructing subspace detectors from the signals generated by multiple events permits the detection of signals displaying far greater variability in signal characteristics. Practically, this should mean increasing the size of the detector's footprint geographically. The tradeoff is that the sensitivity decreases for a required false alarm rate.

In our case study, a primary detector (STA/LTA on beam) was run on the first 3 days of the sequence, resulting in 308 detections. Using the *Builder* program (Figure 3.1.6) it was possible to inspect and quality check the waveforms at the times of all accepted triggers and to delete segments which clearly were contaminated with interfering signals. (A discussion of the effects of intrusive signals is provided in Section 4.5.) Having deleted all such contaminated waveform segments, a high-rank subspace detector was created from the "cleaned" event ensemble. An energy capture requirement of 0.9 resulted in a subspace detector with rank 164. That the rank is so high compared with the number of constituent signals is a clear indication of the waveform diversity of the sequence to be examined. (If all of the master events in the event pool had come from a very compact region, we would anticipate that a lower rank subspace detector would result.)

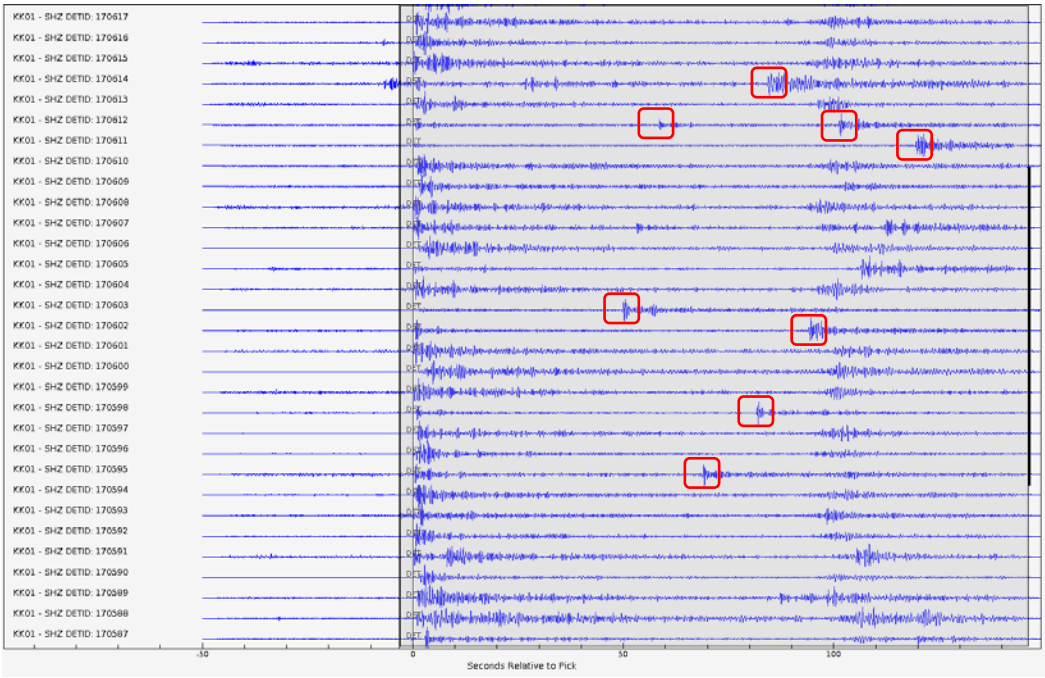

Fig. 3.1.6 Analysis of waveform segments surrounding primary detections in the "Builder" program. Manual inspection identifies instances of templates that, for example, include foreign signals which would degrade the performance of a subspace detector spanning a representative set of events. Examples of these interfering signals are highlighted in red.

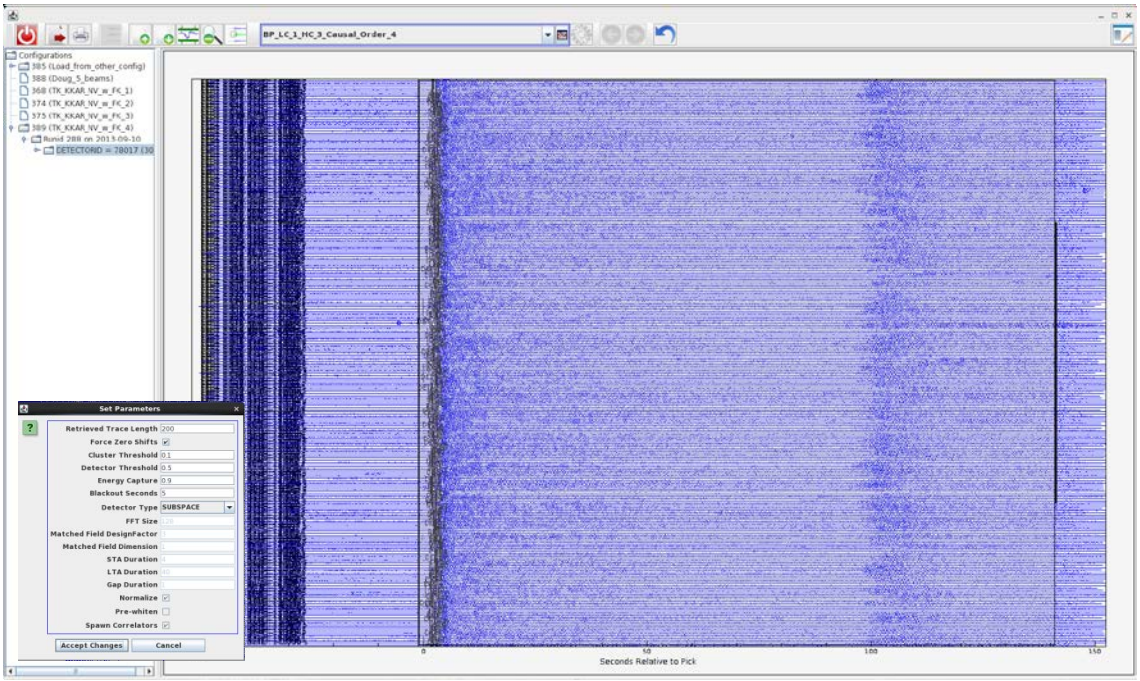

Fig. 3.1.7 Display from the Builder program with the "clean" events used for creating the high-rank subspace detector. The shaded area indicates the time window used (approximately 140 seconds). The data were bandpass filtered between 1 and 3 Hz.

The high-rank subspace detector was run on KKAR data for a period of 21 days, starting 11 days before the main shock. A detection was declared when the subspace detector output exceeded a predefined threshold. The detection threshold used in this case was 0.3. An FK screen with an azimuth tolerance of ±15° was applied to validate subspace detections.

No detections (false alarms) were made during the first 11 days before the main shock and, in the remaining ten days, 70% of the events in the reference bulletin (551/786) were found. As shown in Fig. 3.1.8, these events are distributed over almost the entire earthquake zone, which has a length of about 100 km in the southeast-northwest direction.

An additional 222 events, which were not a part of the reference bulletin, were found by the high-rank subspace detector. These were visually inspected, and none of them were found to be false. The last 50 of the 222 additional events, occurring latest in the aftershock sequence are displayed in Fig. 3.1.9. However, the framework missed as much as 30% of the events in the reference bulletin (235 /786). Fig. 3.1.10 shows missed reference events occurring latest in the aftershock sequence.

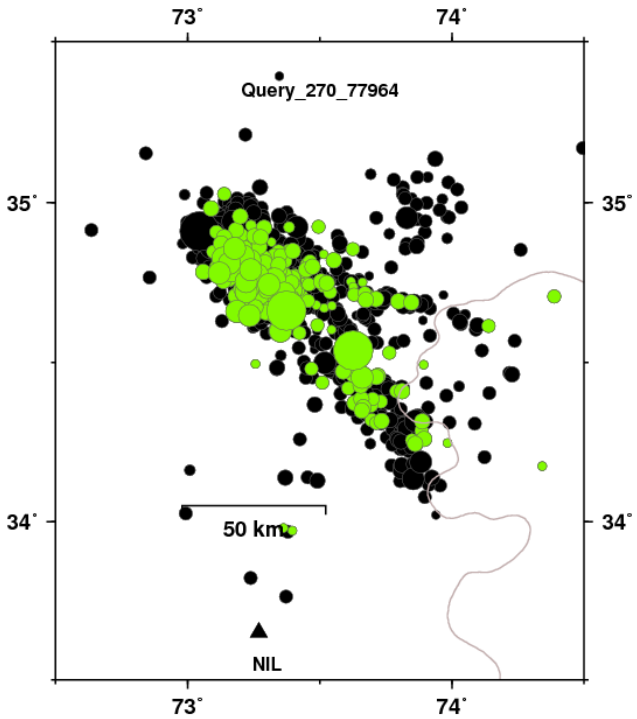

Fig. 3.1.8 Locations of the 551 events in the reference database detected by the high-rank subspace detector on the KKAR array in the time period 2005-270 to 2005-291 (green symbols) and events in the reference database not detected by the subspace detector (black symbols).

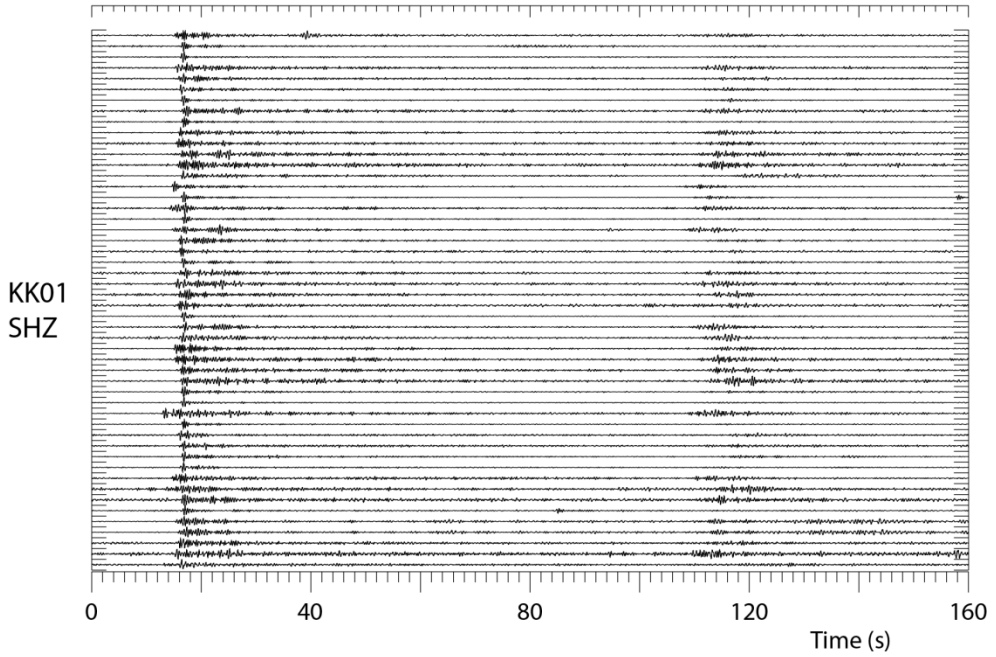

Fig. 3.1.9 KK01 SHZ waveforms of the last 50 of the 222 additional events found by the high-rank subspace detector.

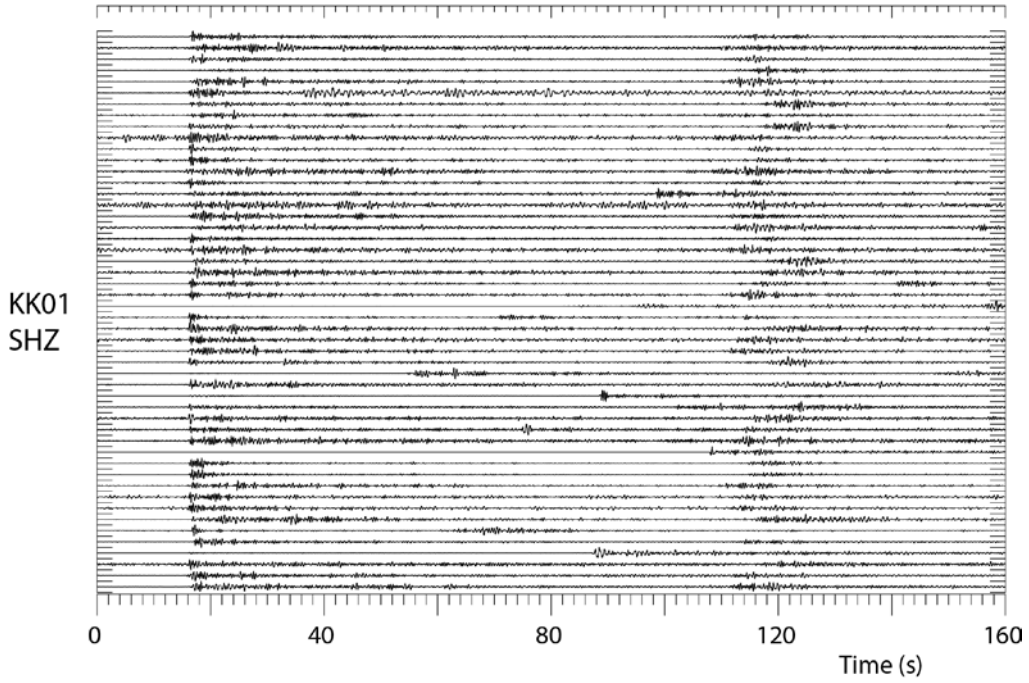

Fig. 3.1.10 KK01 SHZ waveforms of the 50 latest aftershock reference events missed by the high-rank subspace detector.

Since the reference database is based entirely on events found in the ISC and REB event bulletins, no claim is made as to the completeness of the database. It is a truly positive

result that we have found a large number of additional events and that none of these are considered false.

Regarding the relatively large number of missed events (235/786, 30%), there are several factors may have contributed to this. During the relocation effort of the reference events (see Appendix), we were primarily focusing on picking the phase onsets at the different stations. There were no systematic attempts to discard events where the regional distance waveforms contained signals from multiple events. Because of this, a significant number of the missed reference event waveforms contained interfering signals, and these were consequently not triggered by the subspace detector (which was built from "clean" events). Examples of missed reference events with interfering signals are shown in Fig. 3.1.10. This problem of interfering events was most pronounced early in the sequence when the aftershock activity was highest. Section 4.5 demonstrates how a signal can be missed by a correlation detector given a very high amplitude transient in a either of the two correlating data windows and such a cause may not be ruled out for the non-detection of events in the sequence. The flattening of amplitudes in the waveforms recommended by Gibbons *et al*. (2012) may result in different detection characteristics.

Another contributing factor to the missed events can be that the high-rank subspace detector was created form events occurring within only the first 3 days of the aftershock sequence, and may thus not represent the full span of event signatures from the sequence. Lower detection thresholds and less strict screening criteria would also contribute to reducing the number of missed events.

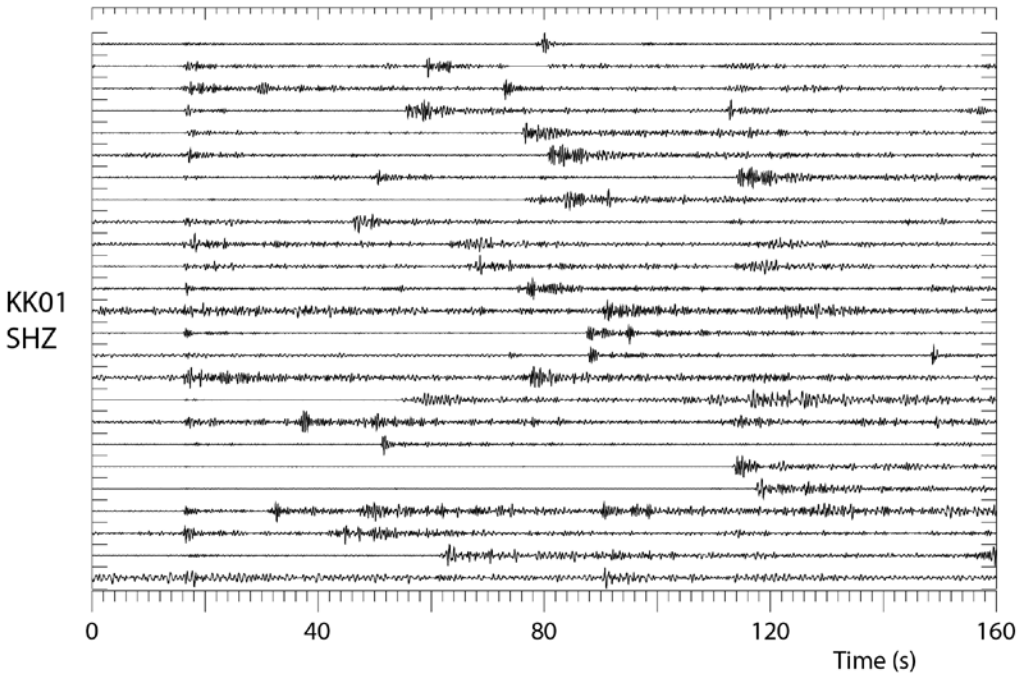

Fig. 3.1.11 Selected KK01 SHZ waveforms of missed events with overlapping aftershocks within the 150 second time window.

The high-rank subspace approach has been demonstrated to be quite efficient in detecting and grouping a very large number of events from the Kashmir aftershock sequence. But, extensive screening (like FK) is necessary to validate the detections. It is therefore not obvious that such an approach would significantly outperform other types of sensitive detectors, when also these are used in combination with extensive screening (like FK).

3.2 The 2012 Sumatra sequence

The aftershock sequence of the 2012 Sumatra earthquake was used to study the performance of subspace detectors to detect and classify events from within a very large (Area = ~250,000 km^2) aftershock sequence observed at teleseismic distances. The ground truth was 1222 aftershock solutions for the first 44 days of the sequence drawn mostly from the Reviewed Event Bulletin (REB).

Subspace detectors were built using several spawning strategies with waveforms recorded by the Makanchi Array (MKAR). Our hope was that subspace detectors could be used to increase the number of detections (relative to a power detector) and to organize the set of detections into groups corresponding to spatially-clustered aftershocks. Ideally, a suitably modified pipeline could use such grouped detections to improve its performance during times of high seismicity. Fig. 3.2.1 shows the geometry and time evolution of the sequence.

We discovered that it was not possible to build detectors using bases of coherent detections, so subspace detectors proved unusable for breaking the sequence into a useful set of grouped events. However, we were able to build high rank subspace detectors that could outperform power detectors and whose detections could with high reliability be assigned to the sequence as a whole.

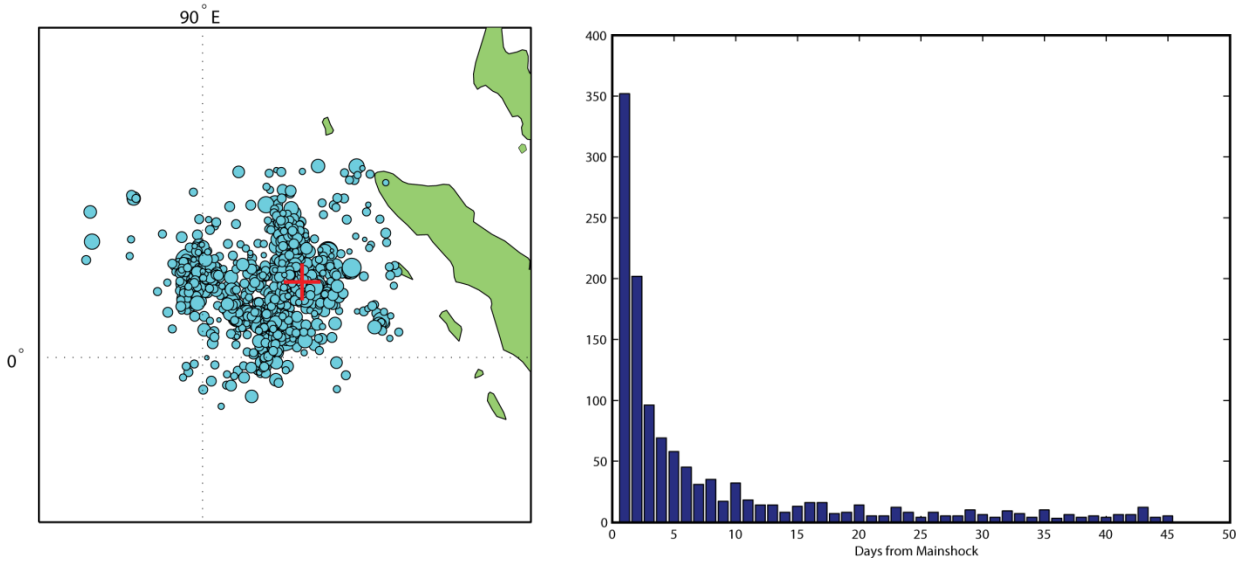

Fig. 3.2.1 Left: The 1222 events from the REB in the source region for 44 days following the main shock. (Right) Events per day for this dataset. Note that by day ten only a few tens of events are being produced in the source region.

We began by performing a series of tests to investigate how the framework would perform on teleseismic data. (Previously, we had only analyzed local and regional data, but MKAR is at a distance of ~ 44°.) These runs also were used to establish the screens applied to power detections used to create templates.

Prior to processing the Sumatra sequence, the framework had not contained any mechanism for screening power detections (or subspace detections) by slowness. We expected to find that FK screens would be very effective at preventing the creation of correlators based on side-lobe detections, and indeed that proved to be the case. For the 44-day interval, one configuration produced 3192 power detections, only 832 of which passed an FK screen appropriate for the region of aftershocks.

A less-expected result was that even when correlators are formed using only power detections that pass the FK screen, it is still possible to get many off-azimuth detections from the correlators. If the SNR of the template signal is too low, then as the detection threshold is decreased the number of false detections increases as well. For example at a detection threshold of 0.1, a correlator whose template had an SNR of ~2.0 had only 0.1% of its detections pass the FK screen. By contrast, at the same detection threshold, a detector whose template had an SNR of ~90 had 59% pass the FK screen.

Based on our preliminary testing, we settled on detection threshold values of 0.3 for the pattern detectors, and an SNR threshold of 5.0 for the spawning detectors. Running the framework with these values produced 366 detections and generated 209 detectors. The top 10 detectors had 131 detections among them and the top detector had 32 detections.

We created multi-rank subspace detectors using detections of the top seven detectors. With this augmented suite of detectors (209 correlators and 7 subspace detectors) the framework produced 480 detections using 216 detectors. Most could be associated to the ground truth events by time and are shown in Fig. 3.2.2. The white filled circles show the event locations from the bulletin. The blue circles represent detections by correlators, and the red circles represent detections from multi-rank subspace detectors.

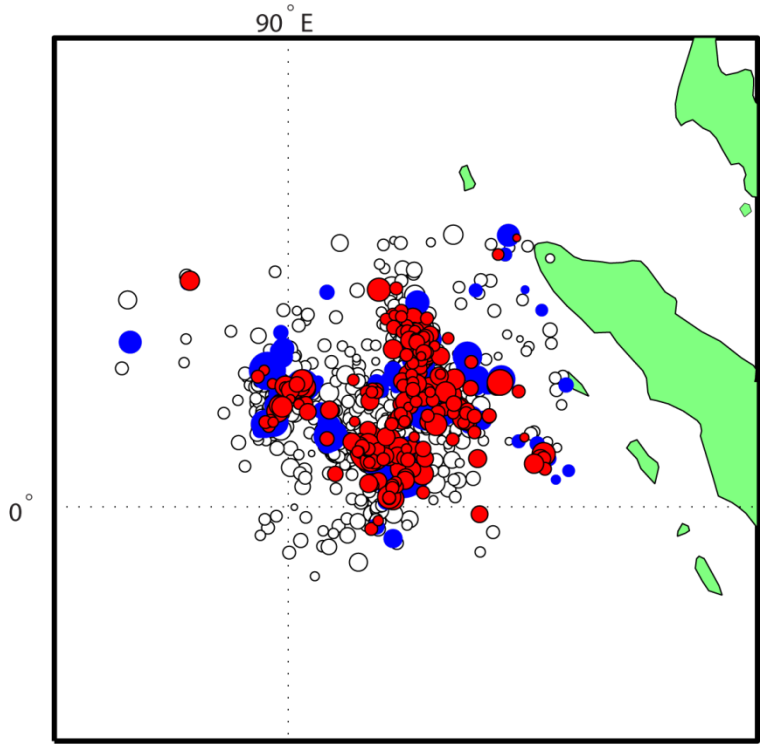

Fig. 3.2.2 (White) Events from bulletin. (Blue) Correlation Detections. (Red) Subspace Detections.

Superficially, the result looks promising. The 7 subspace detectors produced 245 detections and those detections cover a large area. However, the (rank-18) subspace detector with the largest number of detections (129) had an average event spacing of 196 km and a maximum event spacing of 665 km. The second most productive detector (20 detections) had average event spacing of 101 km and a max spacing of 393 km. The centroids of the two sets of detections differ by only 171 km. While some of the scatter is from event mis-location, it seems clear that the subspace detectors are not building spatially compact groups. It is hard to see any useful kind of ensemble processing that can be applied to these groups.

Apparently the basis events are themselves widely scattered despite having been detected by correlation. Fig. 3.2.3 shows one channel of the basis seismograms for the most prolific detector. The combination of narrow-band filtering (1-3Hz) and short duration of the P-pulse produced signals with low time-bandwidth (TBW). Hence the correlations have low significance.

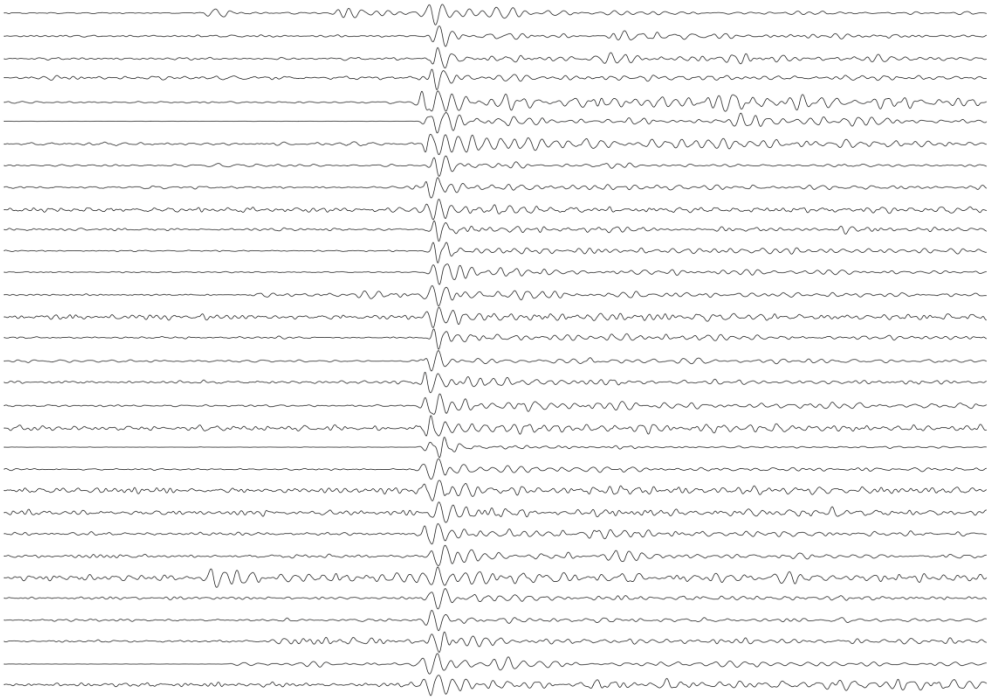

Fig. 3.2.3 Seismograms recorded by array element MK01 for the 32 basis events used to construct the rank-18 subspace detector. Although each of these detections exceeded the correlation threshold, the only commonality of these signals is the short P-pulse followed by low-amplitude coda. The TBW is apparently too low for these correlations to be significant.

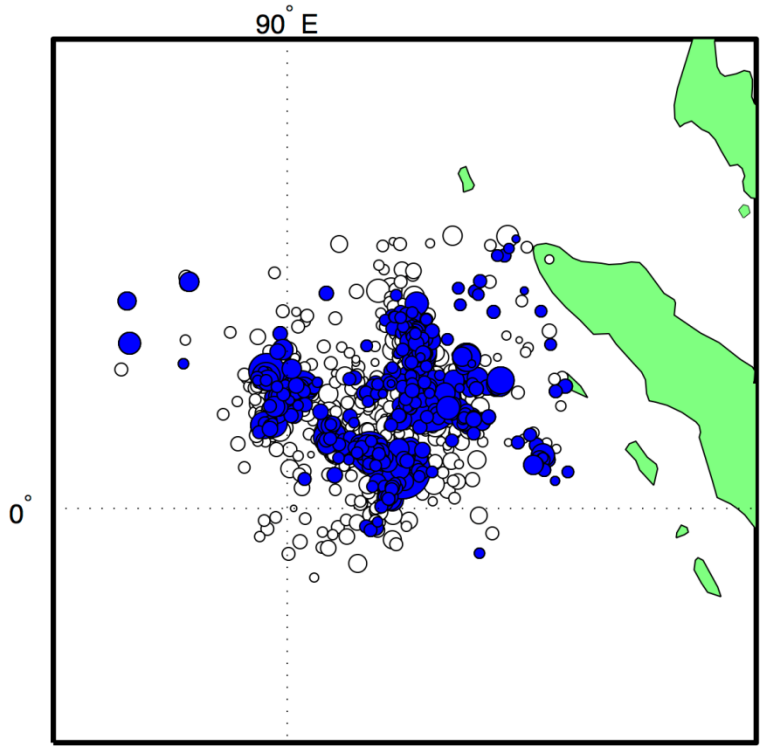

Fig. 3.2.4
485 power detections (blue) superimposed on the 1222 aftershocks from the REB. These were used to create 4 subspace detectors with ranks ranging from 56 to 115.

Although we were unable to use subspace detectors to build spatially compact clusters of detections, we still hoped to demonstrate their utility in increasing the number of detections. The subspace detectors just discussed were not useful for classification, but they were sensitive, having detected far more than just their basis events. By extension, a subspace detector constructed with a basis that spans the entire sequence might be an effective way to maximize the number of detections. Subspace detectors can span the range between pure correlators (rank 1) and energy detectors (rank equal to the signal space). By forming detectors that use more of the signal space in their basis, we can increase the footprint at the expense of an increase in the false alarm rate.

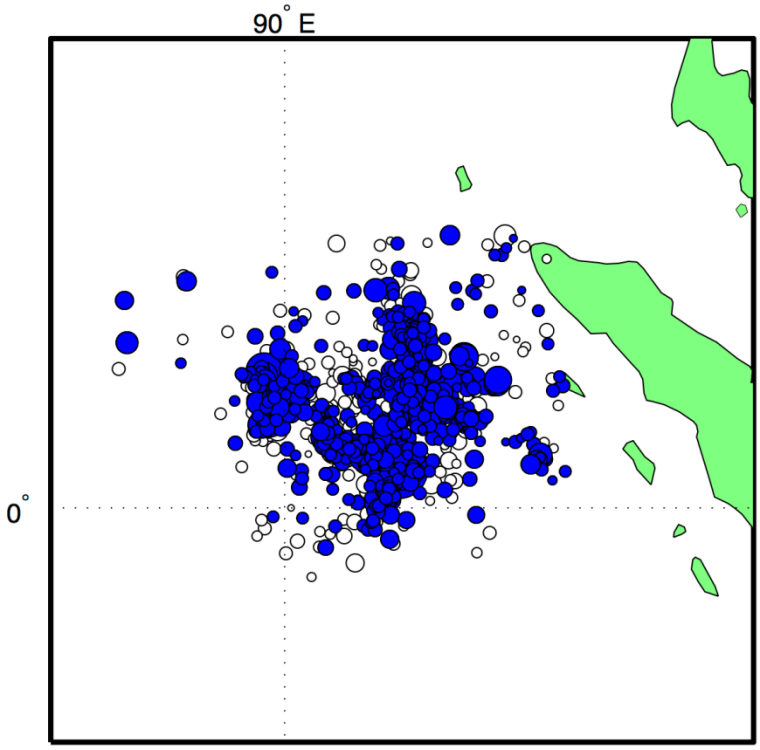

Fig. 3.2.5
629 detections produced by the four high rank subspace detectors and associated to the REB events.

To test this application of subspace detectors we ran four power detectors on the 44 days of data to produce 485 detections. These are shown in Fig. 3.2.4. The four sets of detections were then used to create four subspace detectors with ranks ranging from 56 to 115. Fig. 3.2.5 shows the results obtained running the four high-rank subspace detectors against the 44 days of data. There were 781 subspace detections produced, 629 of which were associated to catalog events by time. In terms of absolute numbers of detections, this is a substantial (63%) improvement over our previous attempt using correlation-derived templates. It seems that although most of the detections are still restricted to the central part of the sequence, within that region, there are significantly fewer missed detections (relative to the catalog) than was the case with the previous

detector. Specifically, within a 4° by 4° box centered on lat = 2°, lon = 92° the catalog contains 889 events while these detectors produced 567 associated detections (63%).

The high-rank subspace detectors produced 152 detections that could not be matched to catalog events, and it is natural to wonder how many of those detections are false. All of the detections passed the FK screen with a minimum FK quality of 0.6, so we can at least say that there was coherent energy at the right slowness to be from the source region of the Sumatra aftershocks. However, for 28 of the detections, visual inspection of a single channel filtered into the frequency band of the detectors, did not reveal any obvious P-arrival. Some of those could be false detections.

We also ran the detectors over ten days of data prior to the main shock. This resulted in 12 detections. Five of those detections could be associated by time to events in the LLNL catalog with epicenters in the aftershock zone. The remainder, although unassociated, are real events with appropriate slowness and good SNR. Fig. 3.2.6 shows these detections on the MK01 channel.

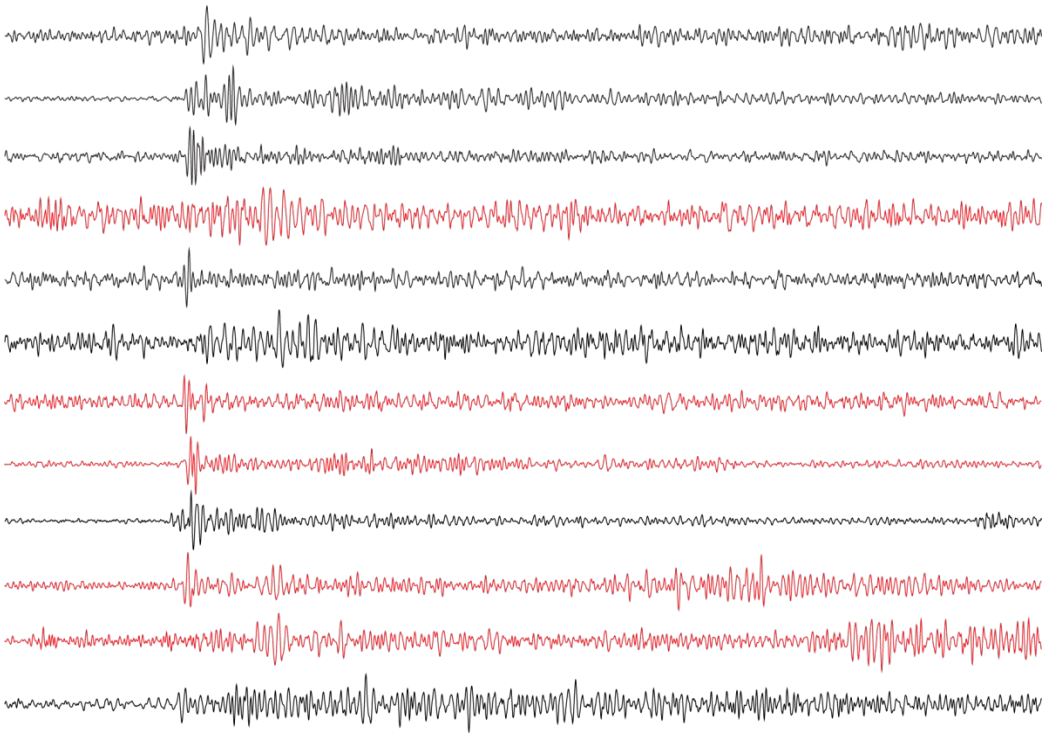

Fig. 3.2.6 All detections (shown on MK01) from the 10 days preceding the main shock. The traces in red are for detections with times that associate with events in LLNL database. The associated events are in the aftershock zone.

These four detectors have a footprint comparable to, or slightly greater than the dimensions of their basis events, and within their footprint, they significantly outperformed the power detectors used to form their basis. However, the implementation used here would not be suitable for a seismic pipeline. In that context one cannot wait until the sequence is over to construct the detectors. Although not discussed here, we have experimented with incremental construction of the subspace detectors, and believe that approach has some promise. However, a bigger problem may be the lack of specificity. As implemented here, high-rank subspace detectors only classify detections as belonging to the overall sequence. While we can imagine ways in which that information might be useful, it would be far better if the classification had a resolution of a few tens of km or less.

3.3 The 2008 Storfjorden sequence

One of the more ambitious studies to date using the detection framework was undertaken by Junek *et al*. (2014, in preparation). This study examined the aftershock sequence associated with the 2008 Storfjorden event with observations from the SPITS array. The objective of the study was to obtain a better interpretation of the underlying tectonic structure driving the Storfjorden sequence. This objective was achieved by constructing, with the framework detections, a vastly larger event catalog complete to magnitude 0.8. The magnitude completeness of this new catalog was fully 1.3 units below that of the best previously-available catalog, the analyst-reviewed NORSAR regional catalog.

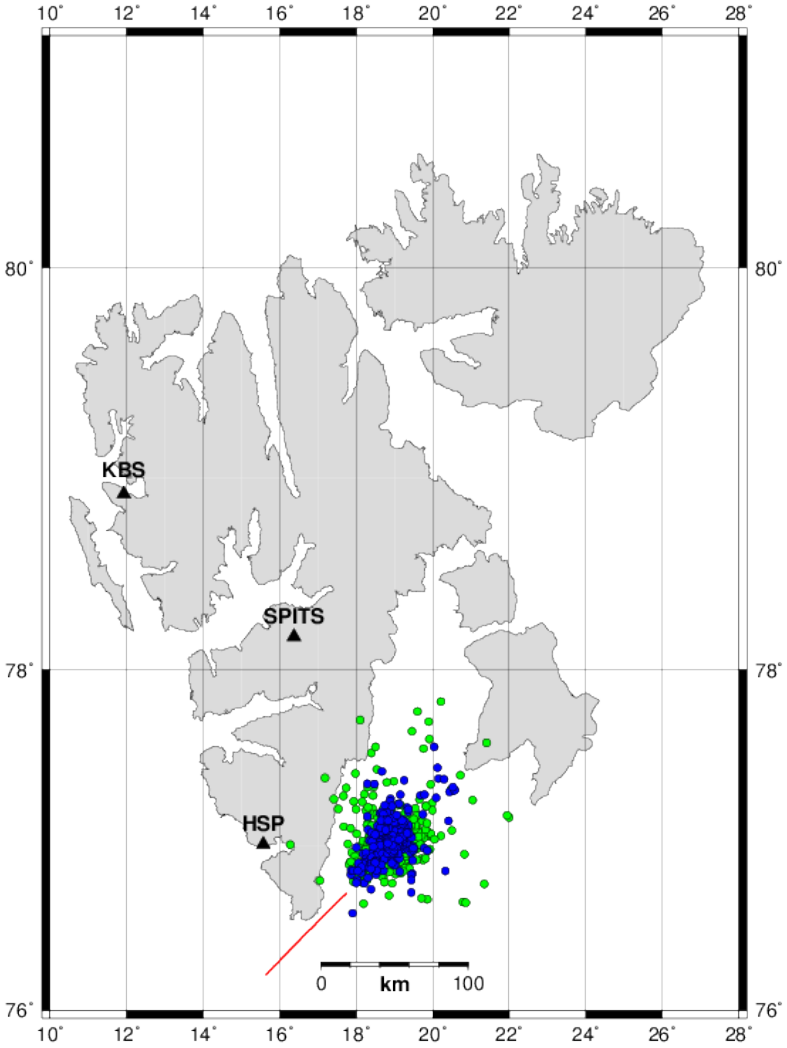

Fig.3.3.1
Map of Spitsbergen and the Storfjorden sequence. The event locations shown in green represent the NORSAR analyst-reviewed catalog between 21 February 2008 and 20 April 2012. Those in blue represent a subsequent relocation catalog (Pirli et al., 2013). The SPITS array is indicated by the triangle approximately 150 km NNW of the sequence.

The Storfjorden sequence was initiated by a magnitude 6.2 earthquake on February 21, 2008 and has produced thousands of aftershocks in the following 6 years. Fig. 3.3.1 shows the position of the sequence off of the southeastern coast of Spitsbergen, which was local to the SPITS array. Because of the difficulties of deploying temporary stations in the arctic environment, no attempt was made to chase aftershocks. Only the permanent stations in the region were used to produce catalogs of the sequence, and these catalogs captured only the larger magnitude events (above magnitude 2).

The detection framework was employed to construct correlation detectors for the sequence over a four-year time window. A single array P beam detector directed at the Storfjorden source region was used as a spawner, and produced over 2000 detectors for the Storfjorden region alone over the four-year period. The detector operated in the 2.0 to 8.0 Hz band and was directed at 150 degrees azimuth with velocity 6.2 km/sec. The framework produced 76688 detections in the four year interval.

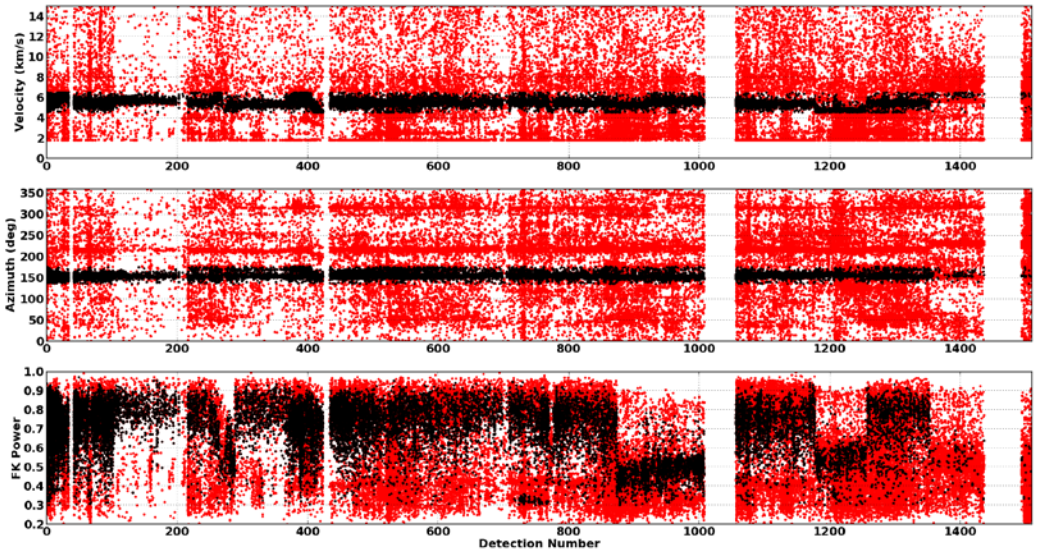

Fig.3.3.2 Scatterplots of detections produced by the framework over a four year period indicate the results of the screening operation taken on detections to limit results to the Storfjorden sequence. The top plot depicts the distribution of events in velocity and time, the middle plot shows the distribution in azimuth and time, and the bottom plot the distribution of an FK quality measure (power in the FK spectrum normalized by power incident on the array). The red dots represent all events and the black dots represent those events passing the FK screen.

Because the azimuth-slowness screen on power detections had not been implemented at the time of this study, the framework produced detectors over the whole region surrounding Spitsbergen, including the nearby North Atlantic ridge and numerous sources producing icequakes on the island. The SPITS array has a small aperture with

relatively high sidelobes. The many sidelobe power detections were responsible for a large number of correlators created for regions other than the Storfjorden source area. Consequently a post-processing step (Figs. 3.3.2 and 3.3.3) was initiated to remove events not from the aftershock source region. An FK screen was used to eliminate events outside the slowness region bounded by 130 and 170 degrees azimuth and 5.2 and 7.2 km/sec. In addition, the screen eliminated detections for which the peak FK power was less than 30% of the total power incident on the array. This post-processing operation reduced the number of detections to 15911.

Following this study, an FK screen was implemented directly in the framework for all power detections prior to the correlator spawning step. Examples like this study have motivated the search for more selective spawning detectors, such as empirical matched field detectors.

Additionally, spot checks of waveforms for the 16,000 events that remained after screening confirmed that S-P times were as expected for the Storfjorden source region.

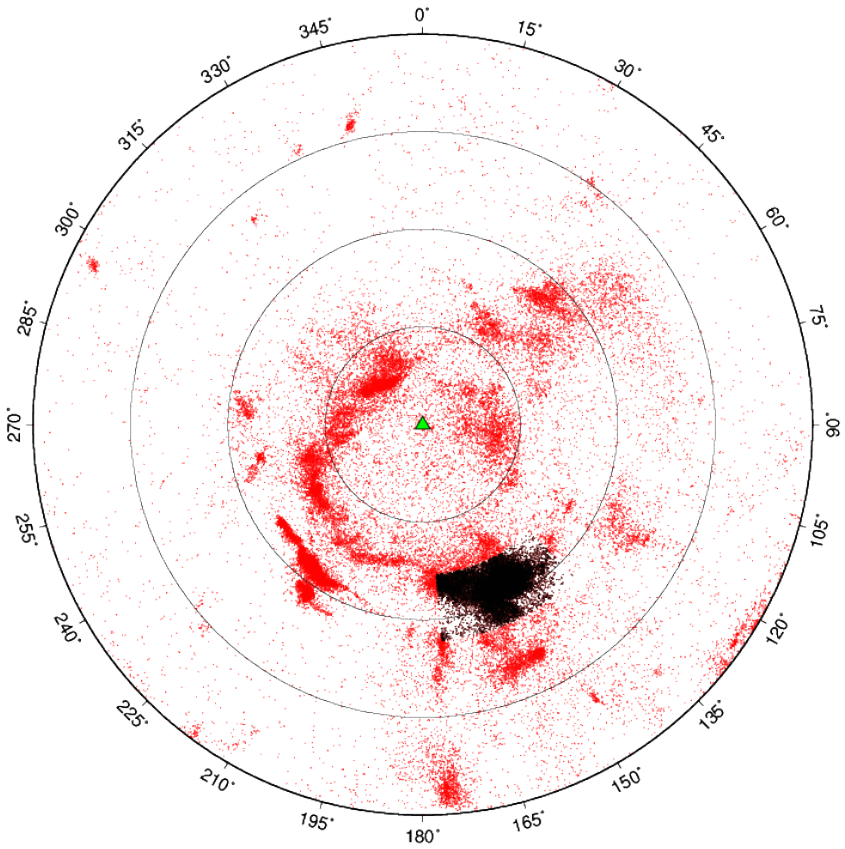

Fig.3.3.3 Distribution of framework detections in slowness space as a polar scatter plot. As in Figure 3.3.2, the red dots represent all framework detections and the black dots represent those detections that passed the FK screen.

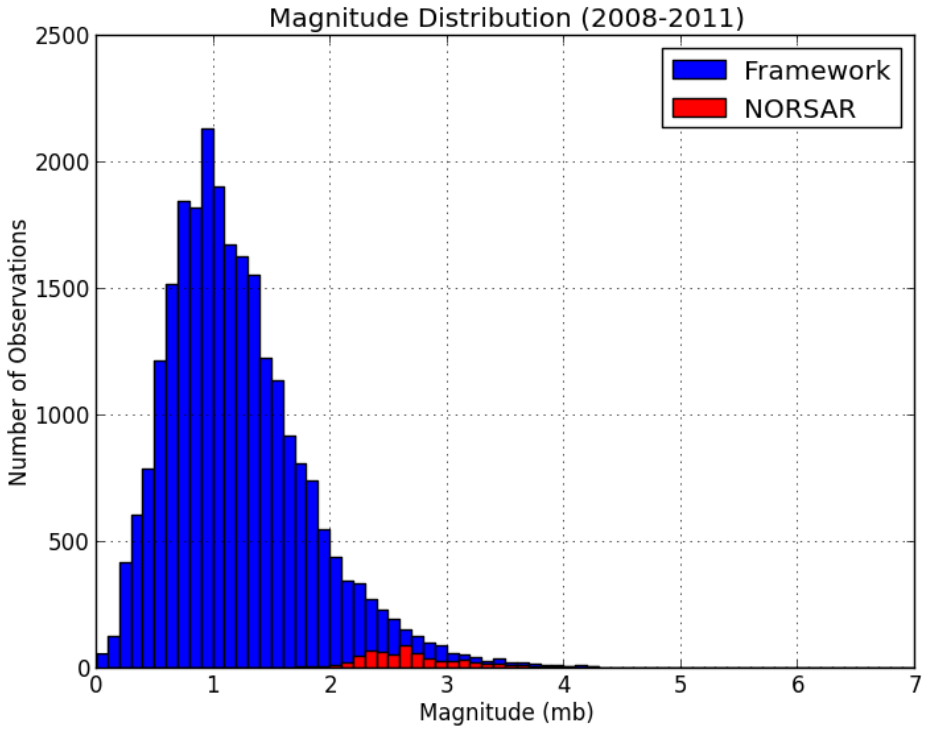

Fig. 3.3.4 Histograms of the magnitude distribution for the NORSAR catalog (red) and the catalog constructed from framework detections (blue) show that correlators produced by the framework collectively have a significantly lower detection threshold than standard pipeline detectors.

From the perspective of this project, the Storfjorden study showed the ability of the framework to reduce detection thresholds with automatically generated correlators and to organize the resulting detections into readily-interpreted groups. Fig. 3.3.4 shows histograms for the NORSAR analyst-reviewed bulletin (red) for the 4-year time period spanned by the study and the catalog of screened detections created by the framework (blue). The framework vastly increased the number of events available for interpretation, especially in the category of events below magnitude 2. Fig. 3.3.5 shows the cumulative frequency-magnitude distributions for the two catalogs and the lines fit to the linear portions of the distributions to estimate the completeness magnitudes. The completeness magnitude for the framework distribution was estimated as 0.8 and that of the analyst reviewed NORSAR bulletin was 2.1. The difference in magnitudes (1.3) represents the reduction in detection threshold afforded by the framework.

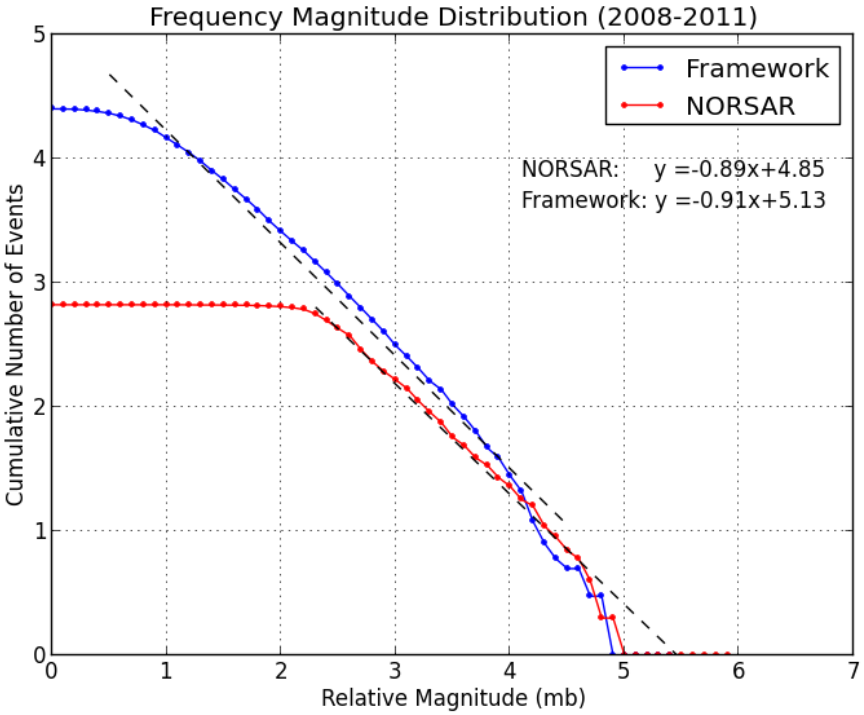

Fig. 3.3.5 Cumulative frequency-magnitude distributions for the framework catalog (blue) and the NORSAR reviewed bulletin (red) indicate a completeness magnitude of 0.8 for the former and 2.1 for the latter. The framework effectively reduced the detection threshold by 1.3 magnitude units.

However, depending on application, a reduced detection threshold alone is not sufficient to recommend use of correlation spawning, as it may significantly increase the analytical work load required to interpret results. It is the ability of pattern detectors to organize detections into groups of closely-related events that provides opportunities for efficient interpretation and discovery of new structure in the data. Fig. 3.3.6 shows the lifespan and size of almost 2500 clusters built by the framework. The cluster lifespan panel in the figure shows the time of creation of a correlation detector (and the cluster to which it corresponds) and the duration of the cluster. The duration is defined as the time of detector inception to the time of the last detection made by the detector.

As can be seen from the figure, cluster lifespans range from a day to nearly the entire sequence history. More than half of the clusters have their inceptions within the first hundred days, and the rest are strung out over the remaining history of the sequence. Cluster formation slows down abruptly around day 80. After this time, the creation of clusters proceeds in fits and starts, indicating an episodic resurgent mechanism for the sequence described by the epidemic model of Ogata (1988). None of this detailed

spatio-temporal structure of the sequence is apparent from examination of the NORSAR analyst reviewed catalog.

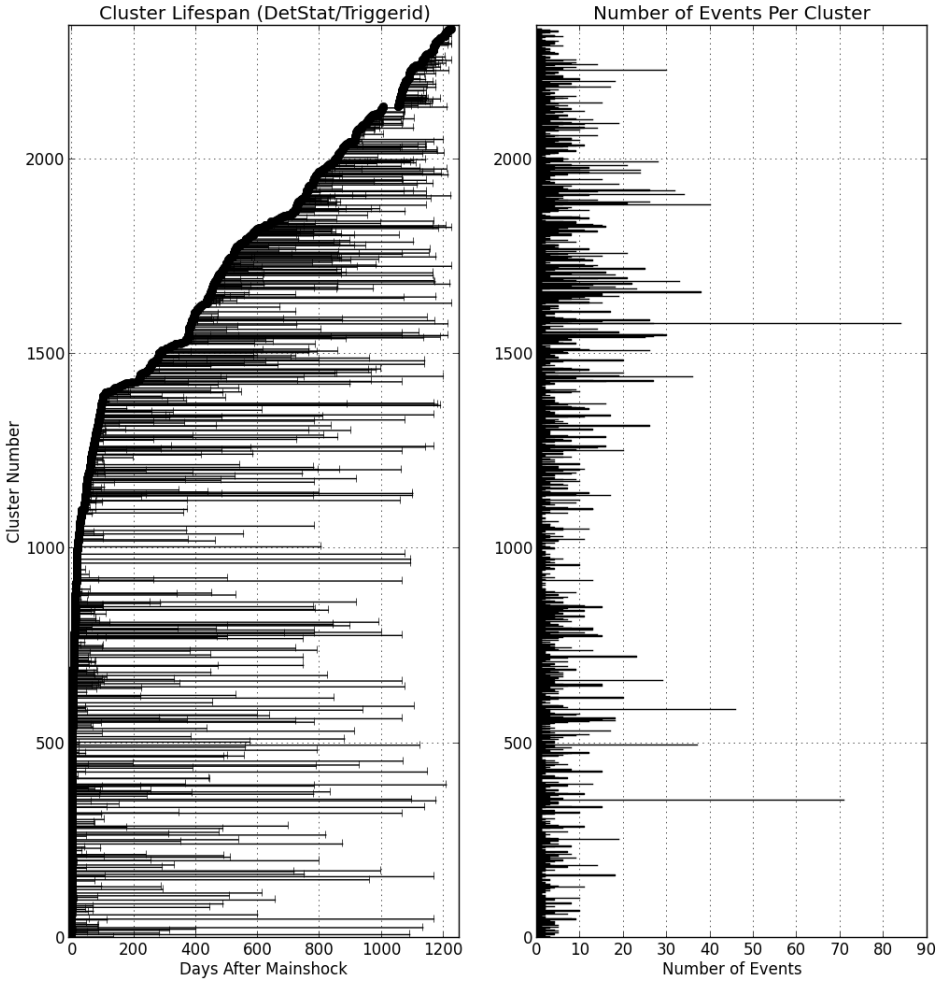

Fig. 3.3.6 Cluster lifespan plot (left panel) and number of events per cluster (right panel) demonstrate the episodic nature of activity in the source region.

3.4 Repeating mining explosions in Western Kazakhstan

The automatic classification of repeating mining explosions for screening has been a long standing issue in the monitoring of nuclear explosions at very low yields. While waveform cross-correlation was identified at a very early stage as being a key method for associating a detected signal with a known source (e.g. Harris, 1991), a detection framework of the form proposed here is necessary for the autonomous implementation of such a classification. Western Kazakhstan is a region of interest for this type of screening given both a large number of repeating industrial sources (primarily open cast mines) and a sparse observational network. Table 3.4.1 provides the locations of mines in the region which have been identified both from local sources of information (Ground Truth information collected by the Kazakhstan National Data Center) and from satellite imagery. The locations of these mines are displayed in relation to the ABKAR seismic array and the AKTO 3-component seismic station in Fig. 3.4.1.

Table 3.4.1 Locations of some known active quarries in Western Kazakhstan and neighboring regions of Russia. Several of these mines are readily identified from Google Earth. Some, for example the Elenovka quarry, are not visible on Google Earth (for which, at the time of writing, the most recent satellite picture is from 2004 for this particular region). The quarry is clearly seen on a more recent satellite photo on the Russian site http://kosmosnimki.ru/

Mine	Latitude	Longitude
Chromtau	50.3231	58.4583
Terensay	51.5605	59.4478
Dombarovskiy1	50.9255	59.4708
Dombarovskiy2	50.8931	59.5300
Elenovka	50.9031	59.8375
Yasniy	50.9817	59.9233
Koktau	50.4808	59.1056
Zhetiqara	52.1400	61.2786

Also displayed in Fig. 3.4.1 are fully automatic single array event location estimates generated only from P and S phase detections on the ABKAR array. Many of these events can be demonstrated to correspond to Ground Truth mine blasts at the sites labelled. While these estimates fall into very clear clusters, the dimensions of the clusters are large relative to the distances between the sites with many events being associated with an uncertainty and bias of the order ~50 km. Additional phase picks from the AKTO station would almost certainly reduce the location error significantly, but the automatic association of these arrivals is non-trivial.

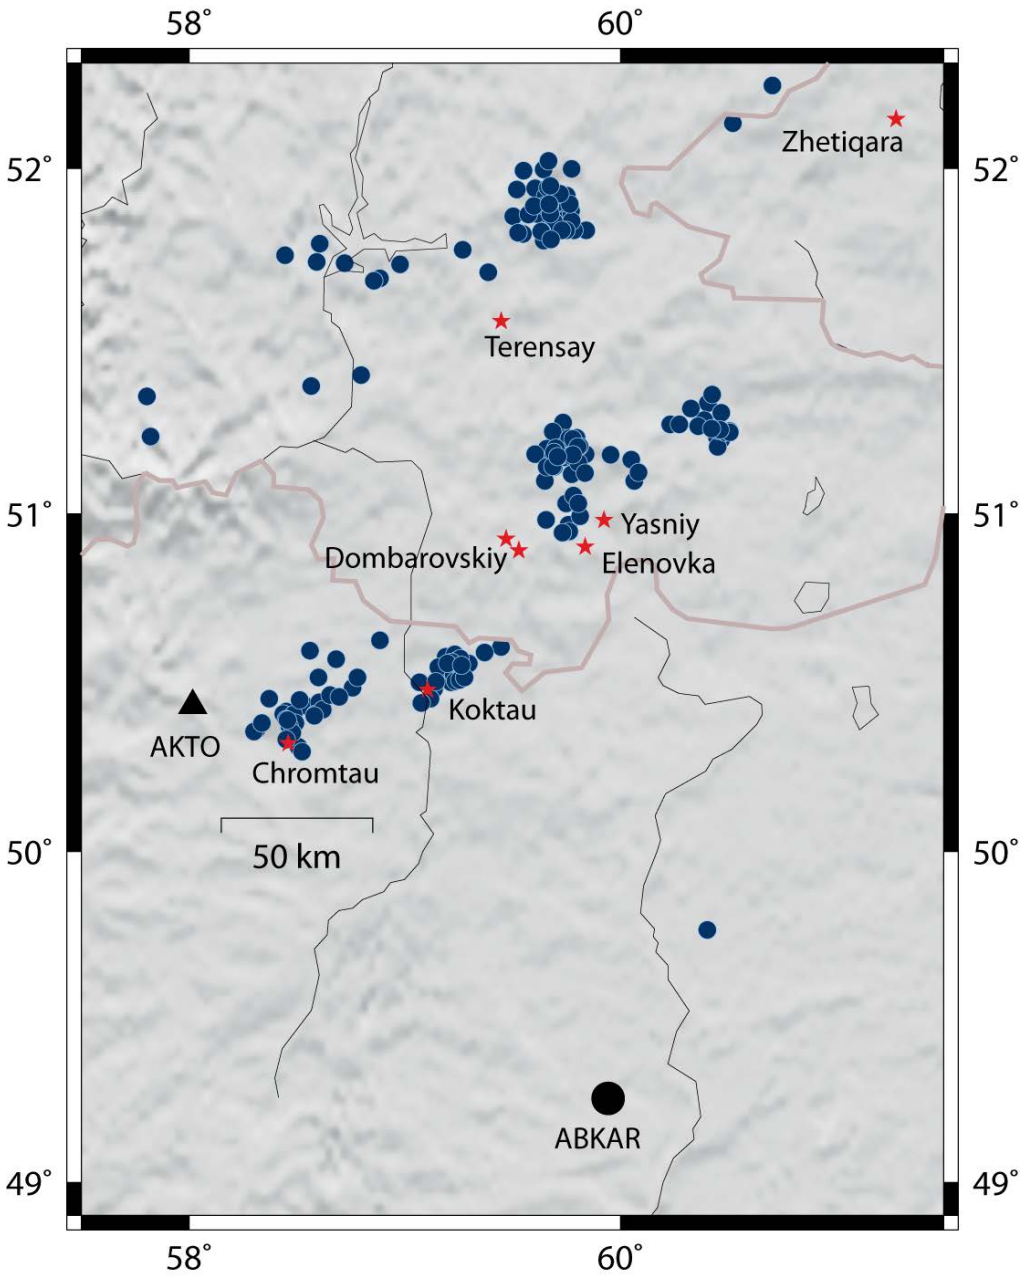

Fig. 3.4.1 Known active mines in Western Kazakhstan and neighboring regions of Russia together with fully automatic single array location estimates of events in the region using only the ABKAR array. The locations of the sites are provided in Table 3.4.1.

The framework was initiated on the ABKAR array with a spawning beam for an initial regional P-arrival with backazimuth 355 degrees and an apparent velocity of 8.0 km/s. An f-k screening was applied such that all detections on this beam that did not indicate a slowness vector for a regional P phase within approximately 10 degrees of the target backazimuth were discarded (see Section 4.4 for details).

Using signal templates of length between 40 and 60 seconds, and the filter band 2-5 Hz, a subspace detection threshold of 0.55 resulted in the clusters as displayed in Fig. 3.4.2 over a period of 90 days. It can be confirmed by analyst examination and event-wise cross-correlation with Ground Truth events from the sites of interest that each of the clusters contains only events from a single one of the sites. That several different clusters arise for each of the sites is a function of the subspace threshold selected and this threshold and the energy capture specification must be adjusted if a different behavior is required.

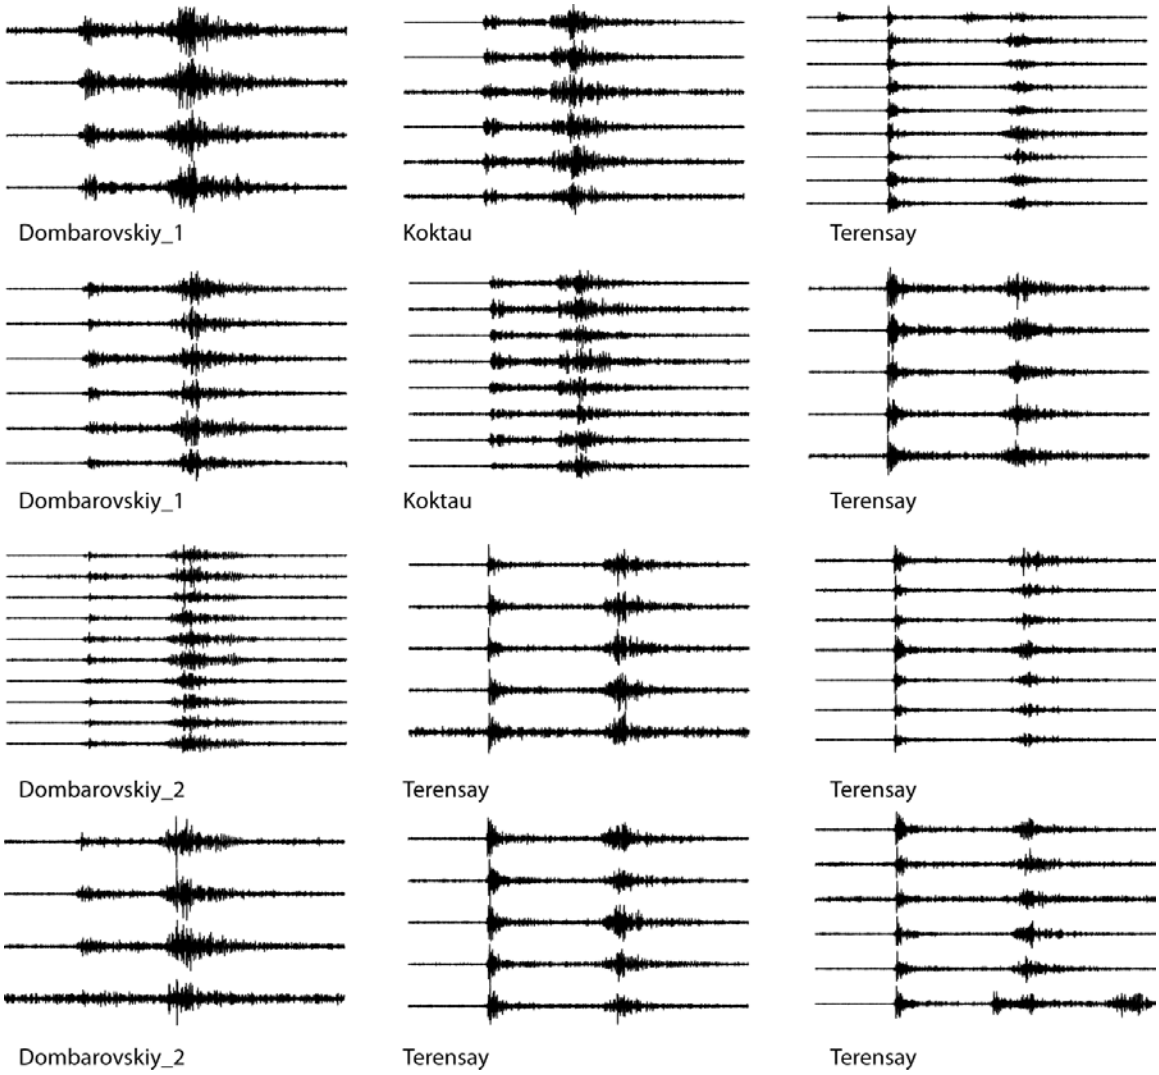

Fig. 3.4.2 The twelve clusters of mining events containing 4 or more signals obtained in the time period 2008-001 to 2008-090 on the ABKAR array. A single spawning beam aimed at backazimuth 355 degrees and apparent velocity 8.0 km/s was used as a primary detector.

4 SPECIAL TOPICS OF INVESTIGATION

In running the detection framework on the case scenarios described in depth in the previous sections, many issues were raised that are generic to the operational implementation of pattern detectors rather than specific to a given monitoring test case. Five such issues are dealt with in the section.

Section 4.1 deals with the scenario of a signal of monitoring interest occurring in an extensive aftershock sequence. We need to ensure that such signals are not erroneously screened.

Section 4.2 addresses the lack of clusters identified in major aftershock sequences using classical correlation-based measures of waveform similarity and investigates, using a coherent multi-array observation of the source region of the 2011 Tohoku earthquake, the idea of using a subspace-based measure of waveform similarity.

Section 4.3 examines the applicability of short-window (single phase) empirical matched field processing as a generator of primary triggers instead of classical STA/LTA detectors on beams.

Section 4.4 addresses the need to perform screening of primary triggers when the target of the detection framework is relatively specific. Since primary detectors are likely to trigger on unrelated signals, depending upon array geometry and signal attributes, we need mechanisms to prevent spawning of pattern detectors for signals not of interest to our monitoring aims.

Finally, section 4.5 examines how correlation and subspace detectors can fail due to the presence of unwanted and unrelated sections of waveforms in specified templates.

4.1 Case scenario of detection of DPRK signals within the Kashmir aftershock sequence

The addition to a standard pipeline of an autonomous mechanism for creating correlation detectors should be at least neutral in the performance of the pipeline, if not actually beneficial. There may be a concern that automatically generated detectors intended to sweep up aftershocks may sweep up desired events as well. In view of this possibility, we conducted a series of experiments to test whether spawned correlators would detect events unrelated to their target aftershock sequence. The specific hypothetical scenario we constructed was the occurrence of DPRK nuclear tests during the Kashmir aftershock sequence.

We used 11 days of continuous data from the ABKAR array for the Kashmir sequence and superimposed (at 1:1 amplitude ratio) the ABKAR observations of the 2009 and 2013 DPRK tests in the stream (Figs. 4.1.1 and 4.1.2). We did not use the 2006 DPRK test in these experiments, since our ABKAR observations of that event lacked one channel. We ran the framework on this composite stream using several configurations, including:

1. A baseline beam recipe similar to what might be used in an existing pipeline (see Fig. 4.1.3). The beams and corresponding power detectors operated in the 1-3 Hz band. The array response pattern is shown in Fig. 4.1.4. The slowness sampling points of the recipe, scaled to wavenumber at the center of the processing band (2 Hz), are superimposed to show that the recipe packs the beams hexagonally at separations equal to one-half of the beam width.

2. The same collection (beam recipe) of array power detectors, augmented with a single array correlator directed toward the DPRK tests as a spotlight detector. This correlator template was constructed from a 65-second window of ABKAR recordings of the 2013 DPRK event beginning about 5 seconds before the P wave. In keeping with spotlight operation a high detection threshold (0.707 linear, 0.5 squared) was assigned to this detector.

3. The same beam recipe, but augmented with a spawning beam and the correlators that it created. In this case, the spawning beam had slowness parameters nearly identical to one of the recipe beams. To avoid conflicts, the closest recipe beam was eliminated from the detector list. No spotlight detector was used in this configuration.

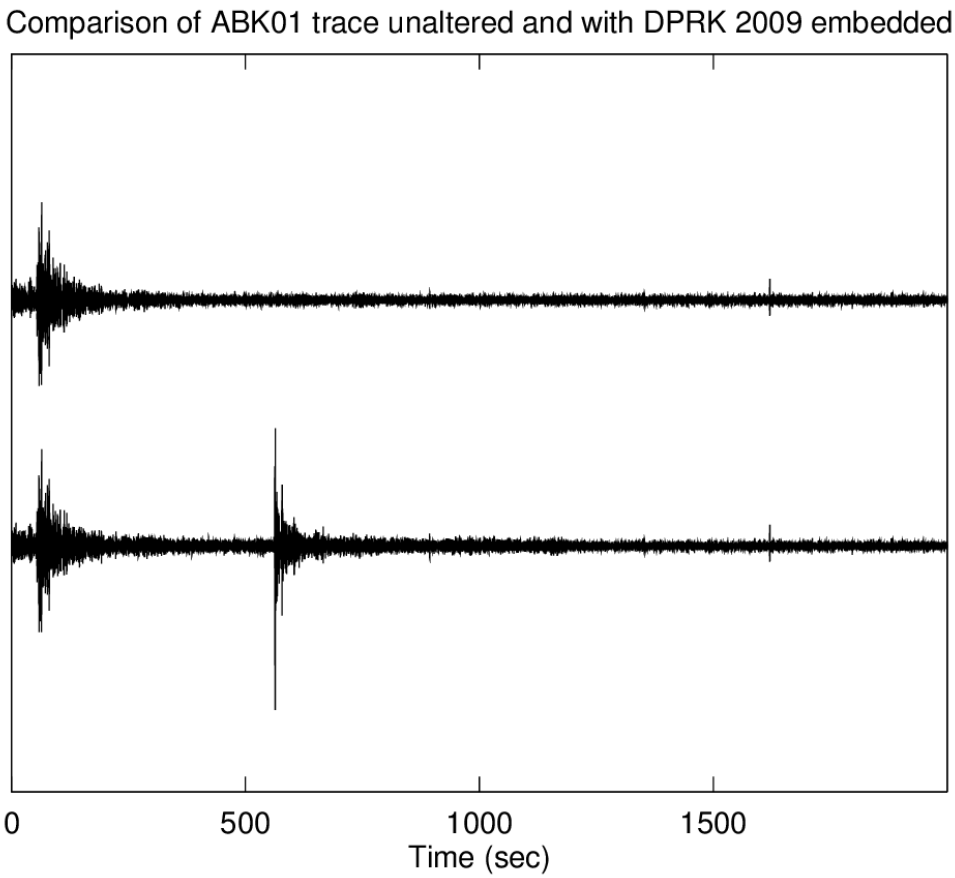

Fig. 4.1.1 Embedding of the ABKAR DPRK 2009 event recording in the ABKAR stream during the Kashmir aftershocks. The embedding ratio was 1:1, so that true relative amplitudes are preserved between observations of the aftershocks and the nuclear tests. One channel (ABK01) before and after embedding and filtered into the 1-3 Hz detection band is displayed. Clearly this embedded observation did not provide a challenging test of detection. However, the point of this exercise was one of classification, not sensitive detection, i.e. to determine whether the correct detector found the event.

The baseline configuration resulted in 266 detections, including the two embedded DPRK events. These two events were detected not by the beam that theoretical slowness and backazimuth considerations would predict (reference Fig. 4.1.3), but by an adjacent beam (90 degrees, 9.24 km/sec), suggesting that unmodeled refraction near the receiver is at play. The distribution of the beam recipe detections is shown in Fig. 4.1.5. In this figure, the symbol color intensity is proportional to the number of detections made by each recipe beam, i.e. the darker the color, the larger the number of detections at that vector slowness. For reference, the asterisk shows the slowness of the spawning detector, which was turned off in this run. Almost half (124) of the detections occurred on the recipe beam closest to the theoretical slowness of the Kashmir aftershock P waves. Several other beams with significant numbers of events appear to

have captured additional aftershocks off the main lobe of the response pattern – this may be an indication of near-receiver refraction for these events.

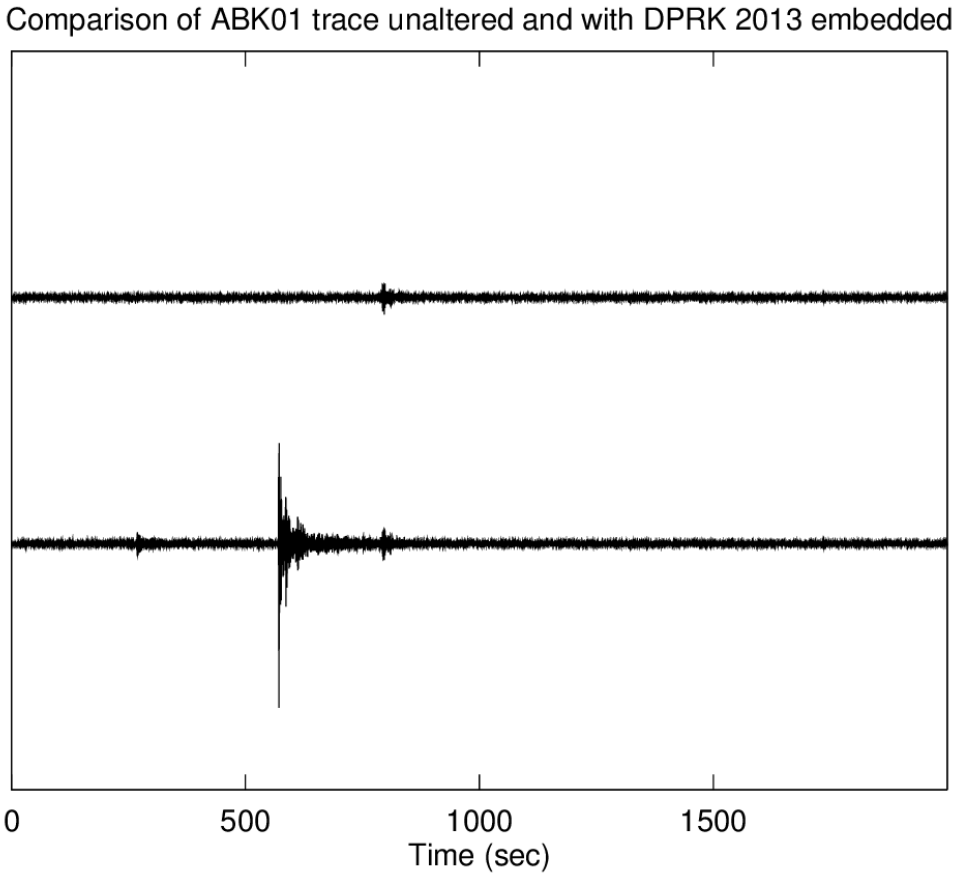

Fig. 4.1.2 Embedding of the DPRK 2013 test in the ABKAR stream about 2/3 of the way through the 11-day Kashmir aftershock test sequence. A second small arrival was recorded about 300 seconds prior to the DPRK P phase in the embedded DPRK segment and was inadvertently introduced into the composite stream. The data have been filtered into the 1-3 Hz detection band in this plot.

The second scenario – the beam recipe augmented by a spotlight correlator – models a likely use of correlators in a traditional pipeline. For this scenario, the framework behaved as one would predict – the two DPRK detections were made by the spotlight correlator rather than the recipe beam detector of the first scenario. Otherwise the distribution of detections was the same as in the first scenario. For this reason we do not show a figure comparable to Fig. 4.1.5 for this scenario. It is interesting that the detection outcome of the correlator is essentially orthogonal to that of the recipe beams for events not fitting the spotlight pattern. As desired, the correlator disturbs the recipe beams minimally.

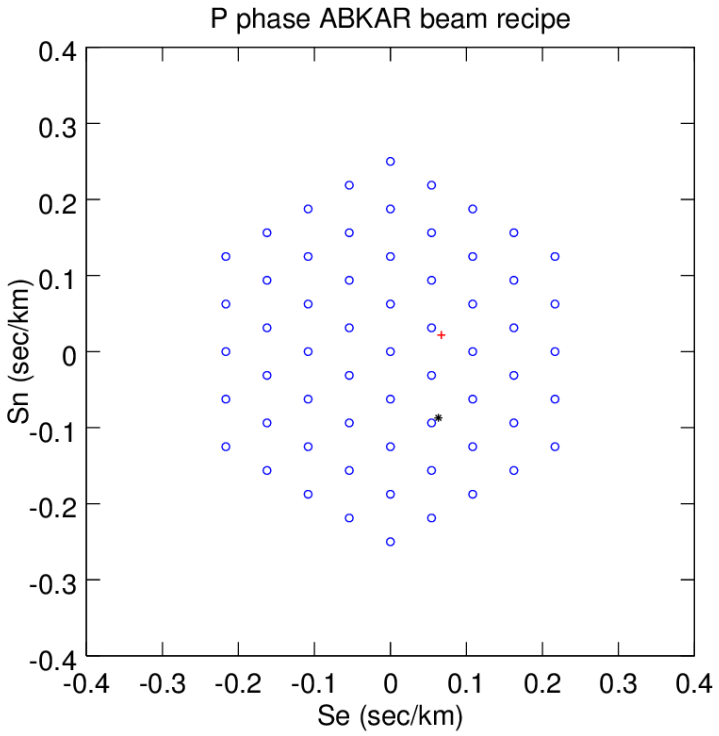

Fig. 4.1.3 P beam recipe for the ABKAR array intended to represent a collection of detectors as might be found in a monitoring pipeline. The beam slownesses are spaced every 0.0625 sec/km, which is about half-width full maximum on the ABKAR array response at 2 Hz (see Fig. 4.1.4). The data were filtered into the 1-3 Hz band prior to beamforming. The red cross indicates the theoretical vector slowness for a P wave originating at the DPRK test site. The closest beam is at 60 degrees, 16.0 km/sec, but the beam that detected the two embedded events was the next closest at 90 degrees, 9.24 km/sec. This result is an indication of P wave refraction near the array. The black asterisk indicates the vector slowness of the spawning beam directed at the Kashmir aftershocks, which replaced the nearest recipe beam.

The third configuration (beam recipe + correlator spawner) resulted in 136 recipe detections and 142 detections by the spawner beam detector and its correlator progeny for a total of 278 detections. The two embedded DPRK events were detected, as in the first scenario, by the recipe beam at 90 degrees back azimuth and 9.24 km/sec velocity. This is the result that we were seeking – the spawned correlators do not sweep up the nuclear tests. Had they done so, it would be an indication that the addition of an autonomous correlator-spawning process to a pipeline would be inadvisable. In fact, for this particular aftershock sequence, the addition of a spawning process did little to alter the outcome of the processing run (compare the distribution of events detected by the recipe beam detectors in this scenario, Fig. 4.1.6, with the baseline distribution, Fig. 4.1.5). The number of event detections was increased from 266 to 278, principally (but

not wholly) by replacing the 124 detections of one recipe beam by 142 spawner and correlator detections. The 142 detections of this latter group comprised 110 detections by the spawning beam and an additional 32 detections by its correlator progeny.

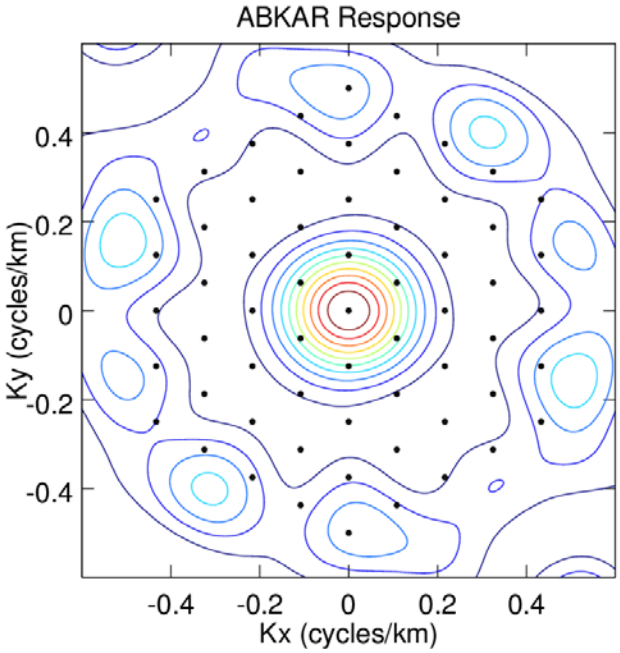

Fig. 4.1.4 Theoretical wavenumber response pattern for the ABKAR array, superimposed with the slowness samples of the beam recipe scaled to wavenumber at 2 Hz.

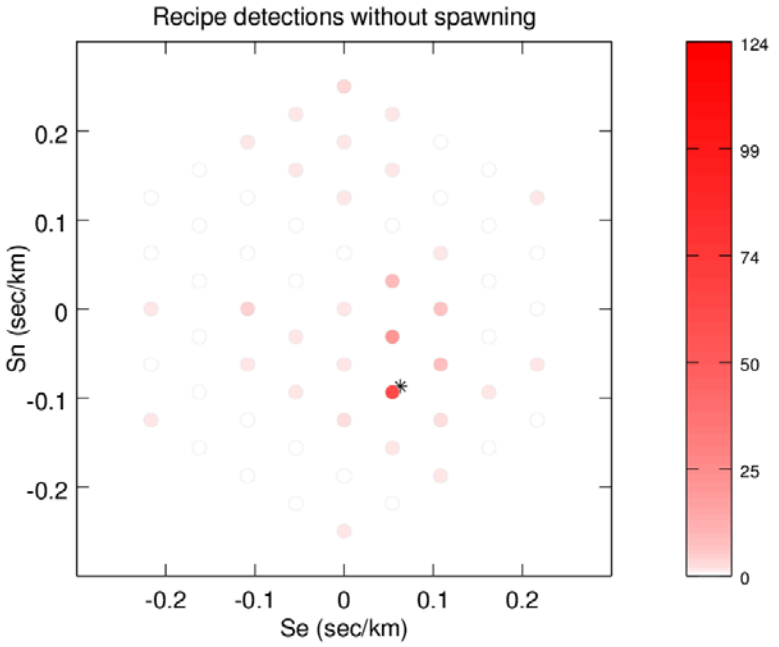

Fig. 4.1.5 Distribution of beam recipe detections by slowness in the baseline calculation. The intensity of color indicates the number of detections by each beam.

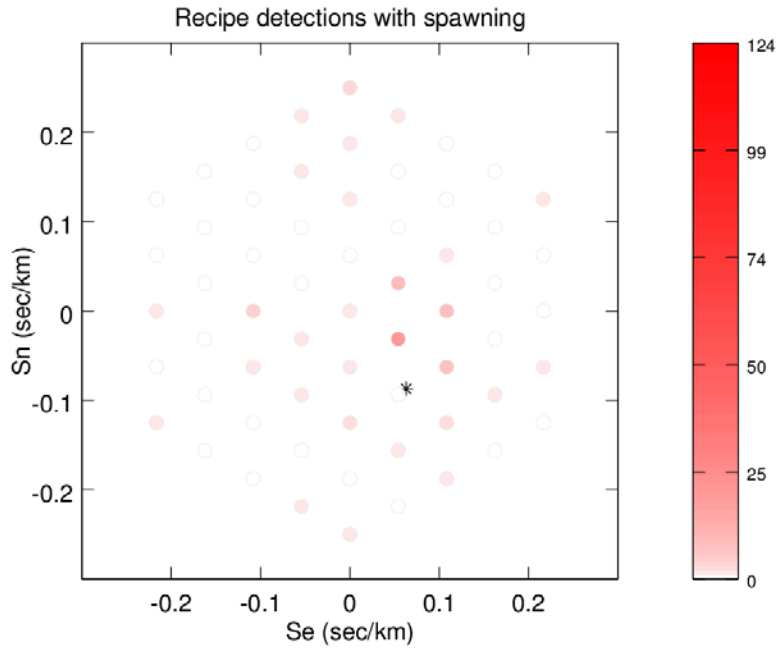

Fig. 4.1.6 The distribution of recipe beam detections with slowness when a correlation spawner directed at the Kashmir aftershocks was added to the framework configuration. The distribution changes very little apart from the substitution of the spawner-derived detections for the 124 detections from the recipe beam replaced by the spawner. For this scenario, the addition of a spawner does little to alter the performance of the detection framework.

One could argue that the total workload for an analyst examining the detections made by the second configuration would be slightly less than the workload required for the first configuration. For the first configuration, all 266 detections would require independent examination, while for the second, the 136 beam recipe detections and the 110 spawning beam detections (246 in total) would require examination. The additional 32 events detected by the spawned correlators would “come for free”, i.e. not require interpretation above that required for the template events (among the 110 spawner detections).

In this view, the addition of a spawning mechanism would cause no harm – perhaps even do a little good, by slightly increasing the number of events detected at a slight decrease in work load. And it would do so without misclassifying the DPRK tests as Kashmir aftershocks.

4.2 Subspace Clustering of the Tohoku Sequence

As reported in section 3, the geographically large aftershock sequences of the 2005 Kashmir and 2012 Sumatra earthquakes produced surprising few clusters using what are now fairly standard agglomerative clustering techniques based on waveform correlation measurements. We carried out a brief attempt to generalize the agglomerative clustering approach to use a subspace based measure of waveform similarity. Our motivation was to ascertain whether a more generalized measure of similarity might lead to subspace detectors that capture and classify more of the seismicity in a large aftershock sequence with a low false alarm rate.

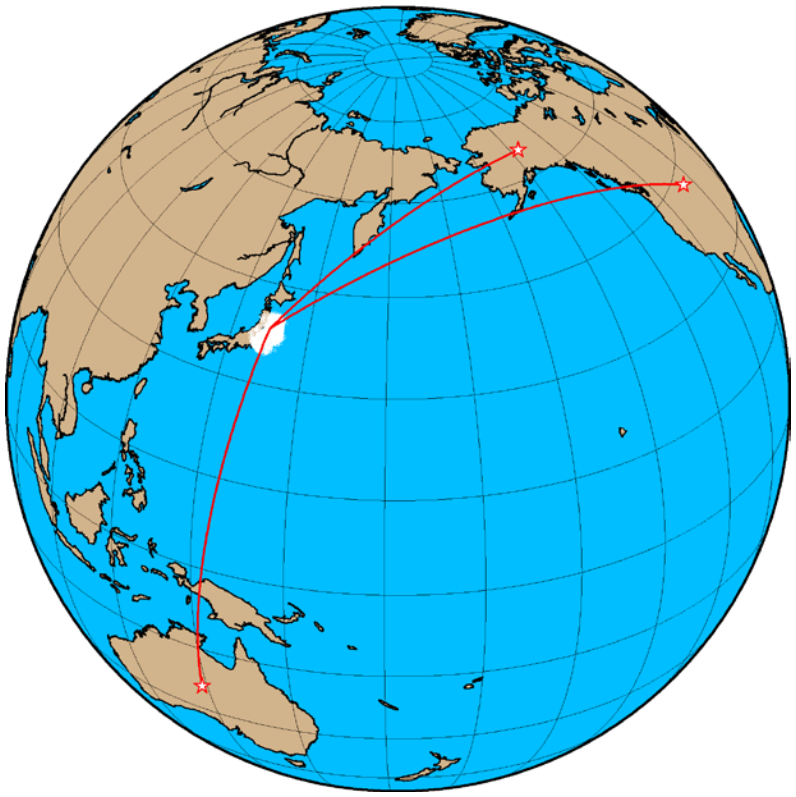

Fig. 4.2.1 Map showing the location of the 2011-2013 Tohoku sequence and the three teleseismic arrays used to observe it: ASAR (Australia), ILAR (Alaska) and PDAR (Wyoming). The cloud of white symbols off the east coast of Japan shows the extent of the aftershock sequence.

As a study example, we examined the aftershock sequence following the 2011 Tohoku earthquake, including events through the end of 2013. A second objective of this study was to test whether coherent combination of waveforms from several arrays in a network would lead to detectors with low false alarm rates. Consequently, we assembled observations of the larger aftershocks of the sequence made by the three

arrays shown in Figure 4.2.1, ASAR, ILAR and PDAR. In the almost 2 year period following the main shock the NEIC reported almost 6000 events of magnitude 4 and greater in the aftershock region shown in Fig. 4.2.1. We downloaded event data from IRIS and manually inspected the 1600 or more events with magnitudes greater than 4.7, retaining 1035 events that appeared to have relatively clean observations on the three arrays (visible signals without overlap between events). These were cross-correlated and the 417 events with correlations above 0.4 were retained for further clustering.

A total of 51 channels of data, being the complete set of vertical channels from the three arrays, were used in the clustering operation.

4.2.1 Subspace Similarity Measurement

The waveforms observed for an event over a network can be thought of as a pattern that characterizes propagation from the event source location to observing stations. A waveform subspace that spans a collection of observations from multiple events in that source region generalizes that notion. Just as two sets of waveforms can be compared with a correlation measurement, two subspaces can be compared with a suitable generalization of the correlation coefficient.

The generalized correlation coefficient is most naturally described in matrix notation. First we define as a vector the collected observations of a single event. If the discrete time waveform observed by channel i, $i = 1, \dots, N-1$ is defined by $x_i[n]$, $n = 0, \dots, L-1$, then the channel-sequential multiplexed form assembles all waveform observations into a single vector $\mathbf{x}$:

$$\mathbf{x} = \begin{bmatrix} x_1[0] \\ x_2[0] \\ \vdots \\ x_N[0] \\ x_1[1] \\ x_2[1] \\ \vdots \\ x_N[1] \\ \vdots \\ x_1[L-1] \\ x_2[L-1] \\ \vdots \\ x_N[L-1] \end{bmatrix} \tag{4.1}$$

If there are M events characterizing propagation from the source region, then the data matrix $\mathbf{X}$ assembles the observations from all events:

$$\mathbf{X} = [\mathbf{x}^1 \quad \mathbf{x}^2 \quad \cdots \quad \mathbf{x}^M] \tag{4.2}$$

An orthonormal basis for these observations is obtained from the singular value decomposition:

$$\mathbf{X} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T \tag{4.3}$$

The basis may have rank d (dimension) up to the number of events, but a lower rank representation is desirable and obtained by partitioning the left-singular vector matrix into a d-dimensional signal-subspace basis $\mathbf{U}_d$ and its orthogonal complement (nominally a basis spanning noise):

$$\mathbf{U} = [\mathbf{U}_d \quad \mathbf{U}_{M-d}] \tag{4.4}$$

To compare subspaces for two clusters of events C_i and C_j, we first compute the projection matrix $\mathbf{C}_{ij}$ between the two representations:

$$\mathbf{C}_{ij} = \left(\mathbf{U}_{d_i}^i\right)^T \mathbf{U}_{d_j}^j \tag{4.5}$$

The rank of the intersection between the two subspaces can be estimated from the singular value decomposition of $\mathbf{C}_{ij}$ (Golub and Van Loan, 1996):

$$\mathbf{C}_{ij} = \mathbf{W}\boldsymbol{\Omega}\mathbf{Y}^T; \qquad \boldsymbol{\Omega} = \begin{bmatrix} \omega_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \omega_{d'} \end{bmatrix}; \qquad d' = \min\left(d_i, d_j\right) \tag{4.6}$$

The singular values of this SVD are one for each dimension (principal direction) that the subspaces share. If the subspaces represented by $\mathbf{U}_{d_i}^i$ and $\mathbf{U}_{d_j}^j$ overlap completely (i.e. the intersection is equal to one or both of the subspaces) then all of the singular values are one. The singular values are cosines of the angles between the principal directions of the subspaces. If the subspaces are orthogonal then all singular values are zero. With these observations in mind, the average sum of squares of the singular values, known as

the normalized affinity (Soltanolkotabi, Elhamifar and Candes, 2013), is a suitable metric of subspace similarity:

$$c_{ij} = \frac{1}{d'}\sqrt{\sum_{k=1}^{d'}(\omega_k)^2} \tag{4.7}$$

This metric ranges between zero and one much like a correlation coefficient. In the event that the subspaces are dimension one as in the case of comparison of two individual events, then the normalized affinity is just the absolute value of the correlation coefficient between the waveforms.

In the subspace clustering algorithm we describe next, it is essential to measure how well a subspace represents the waveforms from a particular event. The metric we use is the normalized energy of projection of the waveforms into the subspace. The projection $\mathbf{x}'$ is defined by:

$$\mathbf{x}' = \mathbf{U}\mathbf{U}^T\mathbf{x} \tag{4.8}$$

The normalized energy (also called energy capture) of the event waveforms is:

$$E = \frac{(\mathbf{x}')^T\mathbf{x}'}{\mathbf{x}^T\mathbf{x}} = \frac{\|\mathbf{U}^T\mathbf{x}\|^2}{\|\mathbf{x}\|^2} \tag{4.9}$$

and simply refers to the fraction of energy of the event waveforms that appears in their projection into the subspace.

To minimize the rank of a subspace representation of the waveforms associated with events in a cluster it is necessary to align the waveforms. This is accomplished by shifting the waveforms of one event as a group (essentially by adjusting the estimate of the origin time) with respect to the waveforms of the remaining events. Thus, an alignment shift is introduced for each event in a cluster to obtain the greatest possibility similarity among the waveform vectors. Referring to equation (4.1), a shift by a single sample in the data is obtained by shifting the rows of the waveform vector up or down by the number of recorded channels N. Zeros fill the vector at the top or the bottom depending upon the sense of the delay.

In practice, the correlation measurement or its subspace generalization (4.7) is maximized over a range of delays pre-defined by constraints on the likely errors of misalignment in the data.

With these definitions in mind, we define the generalization of agglomerative clustering that we attempted on aftershocks of the Tohoku sequence.

4.2.2 Subspace Clustering Algorithm

1. Compute rank-1 similarity measurements (4.7) of all pairs of events. Initially define each event as a distinct cluster C_i. These correlation measurements define a matrix with entries M_{ij}. Only the upper triangular part of the matrix ($i > j$) need be computed.

2. Select the largest correlation measurement M_{ij}. If $M_{ij} < c_{min}$, the termination threshold, then stop.

3. Make a trial assembly $C_{ij} = C_i \cup C_j$ of the events in the two clusters C_i and C_j corresponding to the similarity measurement M_{ij}, aligning the waveforms corresponding to these events in a data matrix $\mathbf{X}$ using the offset defined by the correlation measurement. Compute the SVD (4.3) of the data matrix and retain the column partition of the left singular vector matrix $\mathbf{U}$ corresponding to the top d singular values.

4. Project the data matrix $\mathbf{X}$ into the rank-d subspace (i.e. compute $\mathbf{U}^T\mathbf{X}$) and compute the energy capture of each of the events in the trial assembled cluster C_{ij}. If any of the events has an energy capture that falls below a capture threshold, then reject the trial assembly of step 3 and set to zero the corresponding measurement M_{ij} in the measurement matrix. In that case return to step (2). Otherwise retain the assembly as a new cluster and continue to step (4).

5. Remove rows i, j and columns i, j in the measurement matrix and create a new column and new row; fill these with subspace measurements between C_{ij} and C_k; $k \neq i, j$. Note that this reduces the number of rows and columns in the matrix of measurements by one. Return to step (2).

4.2.3 Clustering Results

The clustering operation described in the last section was applied to the 51 channels of data for the 417 aftershocks known to have correlative twins. This relatively small

subset of events was used in the clustering operation because the computational load is quite large. Even with only 417 events, the repeated calculation of singular value decompositions (steps 3-7) in the subspace similarity measurement for multiple time shifts to align the data is computationally intensive. Another measure taken to reduce computational complexity was to limit the dimension of subspaces to a relatively small value (5 and 9 in two trials). Nevertheless, a parallel code was written for the task and the calculations were undertaken on a 12-core Xeon server.

Another method used to reduce the size of the calculation was to limit the waveform similarity measurement to a 100-second P-phase time window, and to prealign the waveforms from the three arrays by shifting out the respective P propagation times for the main shock location. This operation had the effect of leaving the relatively small differential time delays due to variations in event location from the main shock initiation point.

The clustering operation employed with a clustering threshold of 0.4, an energy capture threshold of 0.4 and a maximum subspace dimension of 5 resulted in many clusters, but particularly ten with at least 5 events. The locations of the events in these ten clusters are shown in Figs. 4.2.2 and 4.2.3. Two of the clusters (Fig. 4.2.3) are geographically compact, but the other eight were diffuse. Four of the clusters had 27 or more events; the largest having 36. The ten clusters had 206 events in total.

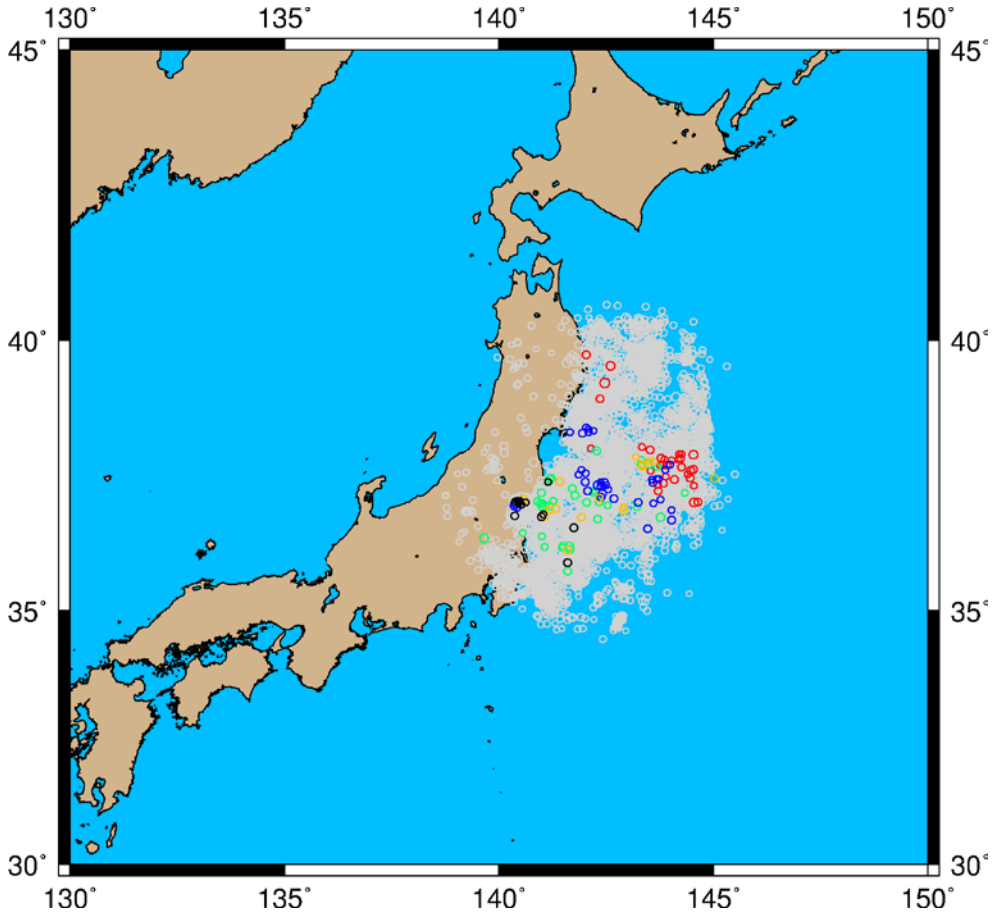

Fig. 4.2.2 Five of the ten largest clusters produced by the subspace clustering algorithm. The white background symbols show the positions of all events in the NEIC catalog. Each cluster is given a unique color. Note that events are widely distributed in these clusters, probably a consequence of the multi-rank nature of the subspace similarity metric.

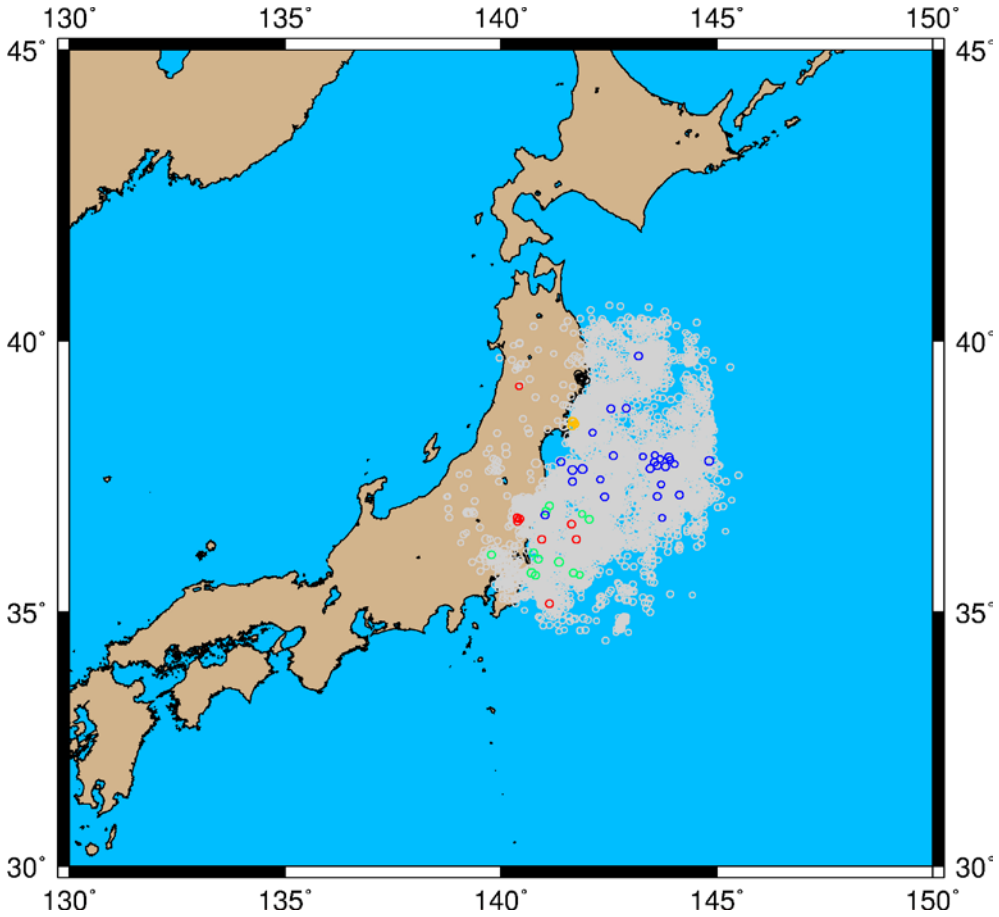

Fig. 4.2.3 The remaining five of the ten largest clusters. Cluster 62 is the compact group of events on the northeast coast of Honshu represented by black circles, and cluster 16 is the large scattered cluster of events denoted by the blue circles.

The two geographically compact clusters (in black and orange in Fig. 4.2.3) produced highly similar waveforms and are consistent with previous results using correlation measures for clustering operations. Waveforms for one of these clusters, highly similar, are shown in Fig. 4.2.4. The waveforms are depicted without the reducing step to align the waveforms relative to the main shock. Hence the bulk propagation delays are evident among the stations. An event cluster of this sort would produce a low-rank subspace detector and would detect relatively few events in the immediate vicinity of the cluster.

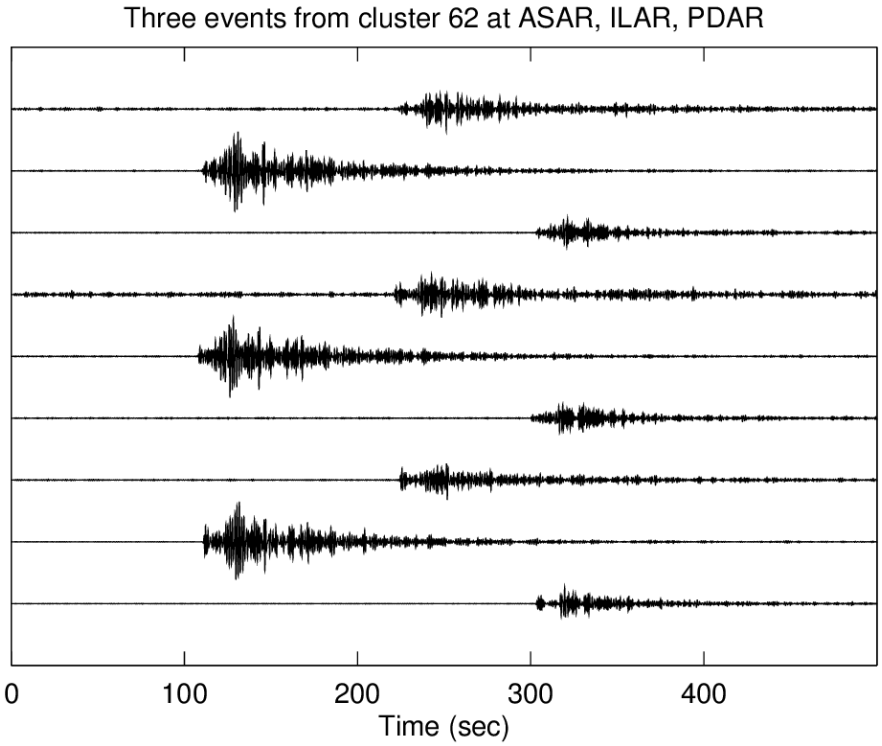

Fig. 4.2.4 Cluster 62, an example of a geographically compact event group. Three events are depicted, with channels AS01, IL01 and PD01 repeated in that order for each of the three events. Note the high similarity of the waveforms. This cluster is indicated by black symbols in Figure 4.2.3.

Fig. 4.2.5 shows waveforms from three events drawn from the very diffuse cluster shown in blue in Fig. 4.2.3. As would be anticipated, there is quite a bit more variation in the waveforms among these events than in the geographically compact cluster, noticeably in relative timing among P arrivals at the arrays. This result is possible only because of the subspace measure being used in the clustering operation. It is surprising that these events were aggregated by an algorithm based on waveform similarity measurements, but the energy capture check ensures that the single rank-5 waveform basis representing the cluster captures 40% of the energy in each of the member event waveforms. It might be anticipated that that a rank-5 subspace detector built from these events would detect a larger number of more diverse events in the aftershock sequence. The hope is that, because the detector would be low rank and have a high time-bandwidth product, it would not make detections outside of the aftershock sequence (which would be considered false alarms). The analysis of this sequence is not yet complete, lacking an actual run of the detectors implied by these clusters against the actual data stream.

A more desirable outcome would be clusters that are geographically somewhere intermediate between the very compact cluster of Fig. 4.2.4 and the diffuse cluster of Fig. 4.2.5. The question of how to achieve this outcome with an autonomous clustering algorithm operating just from the waveform observations is a topic for future investigation.

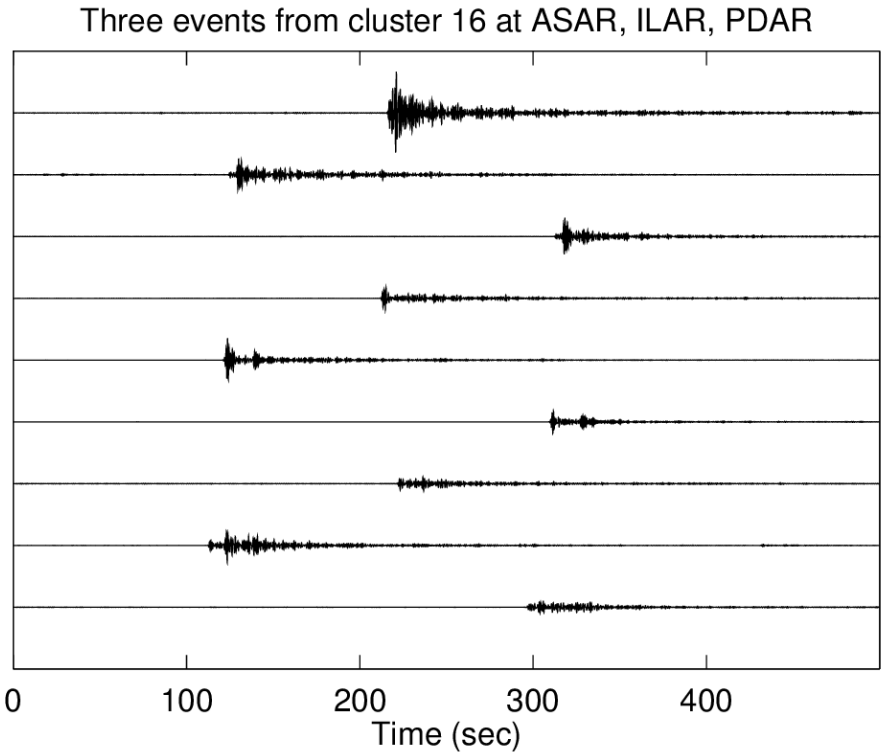

Fig. 4.2.5 Waveforms of three events drawn from a geographically diffuse cluster, the group indicated by blue symbols in Fig. 4.2.3. These waveforms are much more variable, most noticeably at ASAR, the first station of each triple of waveforms.

4.3 Investigation of single-phase empirical matched field processing as primary detectors

The primary detectors discussed so far have been STA/LTA detectors on plane-wave beams steered towards the anticipated direction of the first incoming wavefront from the source region. Empirical matched field processing (EMFP) is a narrow frequency band procedure with origins in underwater acoustics, which has shown great promise as a signal classifier in seismology. Whereas a correlator compares the temporal pattern of a signal from one event with the temporal pattern of signal from a subsequent event, EMFP compares the pattern of phase and amplitude shifts in each of many narrow frequency bands between the different sensors of an array. Harris and Kværna (2010) demonstrated that EMFP was a far more effective classifier of the seismic signals generated by mining blasts than a straightforward correlator. The source-time function of each of the ripple-fired blasts was very different, resulting in very different temporal seismic signatures, whereas the narrow-band phase structure of the wavefront over the array was highly characteristic for each source region. The template has the form of a set of *empirical steering vectors*, calculated from the spectral covariance matrix of the data for reference events for each source of interest.

EMFP is essentially a generalization of a plane-wave beamformer. In the latter, the phase shifts applied to each channel at each frequency are specified by a plane-wave model. In EMFP, these phase shifts are calculated directly from the signals generated by previously observed events at the site of interest (hence empirical) and so can correct implicitly for deviations from the plane-wave model resulting from scattering and diffraction along the path. In the next few plots, we present evidence that EMFP may provide a highly sensitive and source-region specific primary detector for the spawning of new events, which is far less likely to trigger on unrelated signals and so reduces the need for screening.

Fig. 4.3.1 displays the correlation between the signals on the KKAR array generated by the M=7.6 main event and an M=5.5 aftershock from the October 8, 2005, Kashmir sequence. It is clear that the signals have very different temporal structures, primarily due to:

- differences in magnitude (and therefore rupture dimensions)
- different geographical location (the source region of the Kashmir aftershock sequence has an aperture of approximately 80 km)
- differing spectral content.

The clear waveform dissimilarity between these two events rules out the possibility of classifying the entire aftershock sequence using a simple correlator.

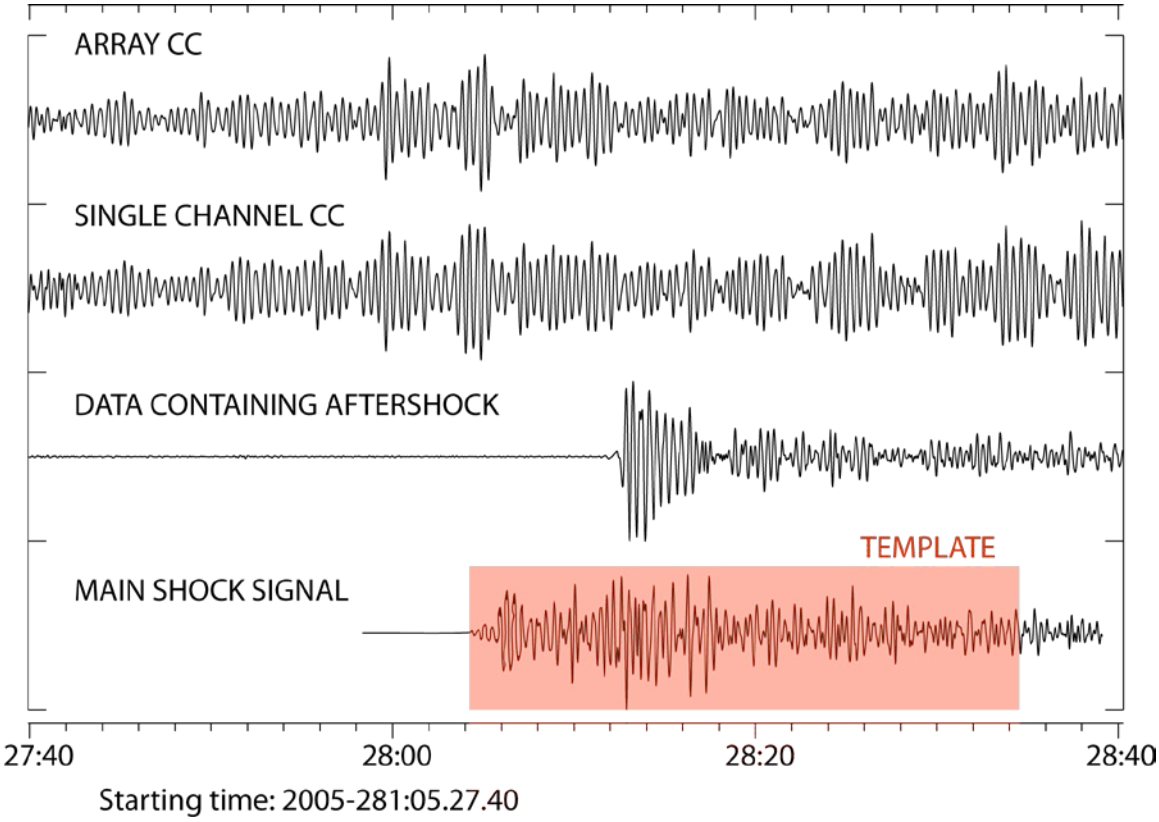

Fig. 4.3.1 Illustration of cross-correlating the signal on the KKAR array generated by the Kashmir main shock (starting time of 30 second long template: 2005-281:03.52.50.385) with the signal from a significant aftershock some 90 minutes later. Both of the signals have high SNR but clearly have very different forms and do not result in correlation coefficients significantly above the background level. It is, for example, not possible to measure the time-delay between signals using a straightforward cross-correlation, even using the array stack.

At a distant seismic array (KKAR is approximately 9 degrees from the sequence, by far the closest observing seismic array) the wavefronts from each event are likely to approach the station from very similar directions. At these distances, even events which are separated by several tens of kilometers are likely to generate wavefronts with very similar projections onto the horizontal slowness space at the array.

Fig. 4.3.2 shows the narrowband f-k spectra (using multitaper estimation of the spectral covariance matrices) for the two arrivals shown in Fig. 4.3.1. On the left hand side are the slowness grids at 2 Hz for the main shock signal (top) and aftershock signal (bottom) and on the right are the corresponding panels at 4 Hz. Note first that there is a

qualitative difference between the peaks in the f-k spectra between 2 and 4 Hz for both events. Whereas the backazimuth indicated for all panels is close to the great-circle backazimuth of ~155 degrees, the apparent velocity measured at 4 Hz is significantly higher than that measured at 2 Hz and, without additional information, is likely to be interpreted as coming from a more distant event. The frequency bands between 2 and 4 Hz display a gradual increase in the apparent velocity estimated and, above 4 Hz, the coherence between the sensors of the KKAR array is reduced to the extent that no significant slowness vector is indicated using classical plane-wave f-k analysis. A narrow frequency band analysis as displayed in Fig. 4.3.2 provides an intuitive picture as to why plane-wave f-k estimates can display such high variability as a function of frequency (see, e.g. Gibbons *et al*., 2010).

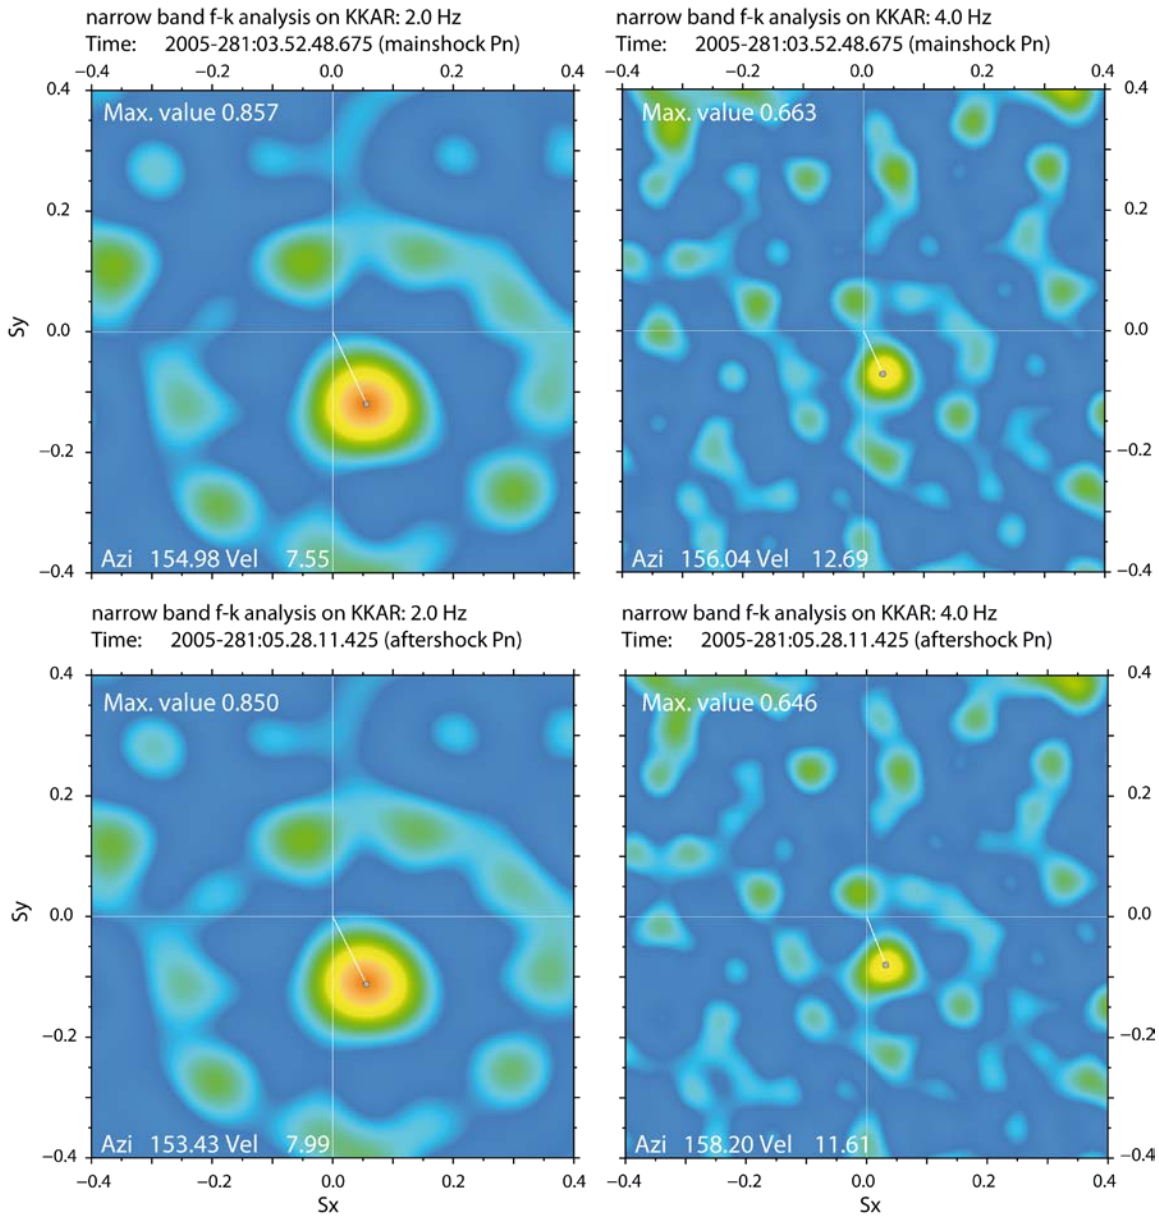

Fig. 4.3.2 Left panels: narrow band f-k analysis when filtering with a center frequency of 2 Hz. In the right panels, the center frequency is 4 Hz. In the top panels we used data from the time segment which corresponds to P_n phase arrivals from the main event. In the bottom panels, a time segment associated with a significant aftershock was selected.

Significantly, despite the dissimilarity between the temporal forms of the two wavefronts, the projections into slowness space for the two narrow frequency bands are remarkably similar for the two events. If we calculate an empirical steering vector for each narrow frequency band for the main shock first arrival, and then calculate the narrowband f-k spectra at the time of the aftershock first arrival - relative to this empirical steering vector, we obtain the relative f-k spectra displayed in Fig. 4.3.3.

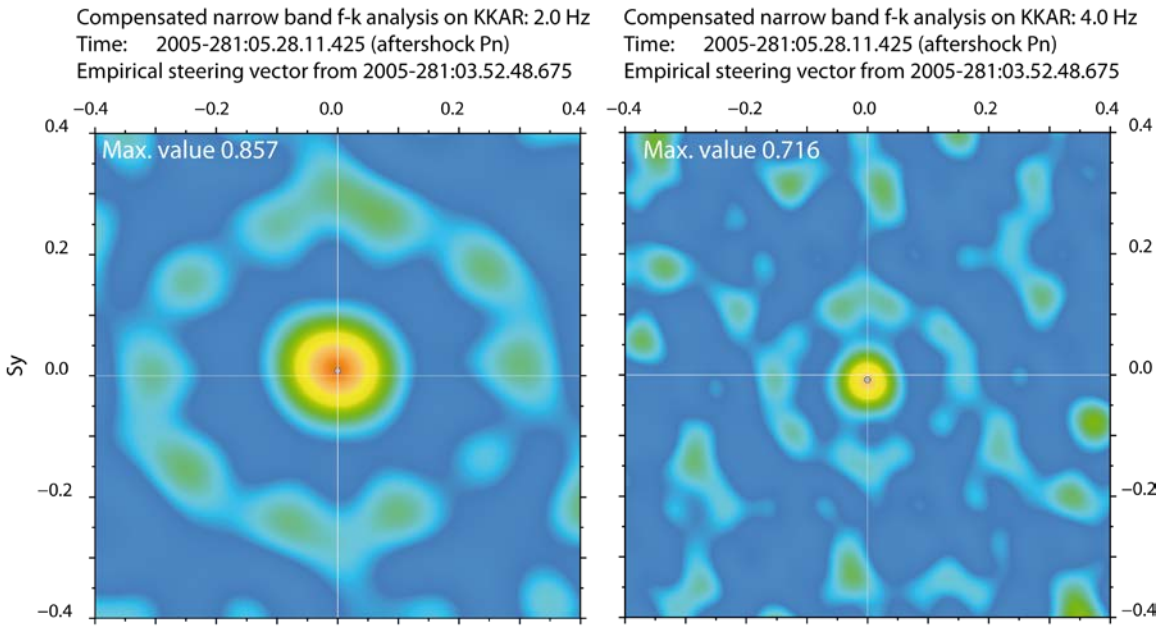

Fig. 4.3.3 Narrowband f-k grids for the first arrival of the aftershock generated after pre-steering by the empirical matched field steering vectors calculated from the main shock first arrival at 2 Hz (left) and 4 Hz (right).

The relative matched field f-k spectra displayed in Fig. 4.3.3 now measure the deviation from the direction of arrival (DOA) of the template steering vector, rather than the DOA of the incoming wavefront itself, and are both almost centered on the zero slowness vector. The value at the center of the grid, the matched field statistic or relative power, is higher than the relative power for the f-k spectrum without the empirical steering vector (lower two panels of Fig. 4.3.2). The gain is marginal for 2 Hz (0.857 as opposed to 0.850) but more significant at 4 Hz (0.716 as opposed to 0.646) since the higher frequencies correspond to shorter wavelengths of ground motion which are more susceptible to perturbation by scattering and diffraction. The centering of the peaks is a useful property since we can measure the validity of a detection based only upon the distance of the peak from the center of the grid, rather than the distance from a frequency-dependent reference point for which calibration is required.

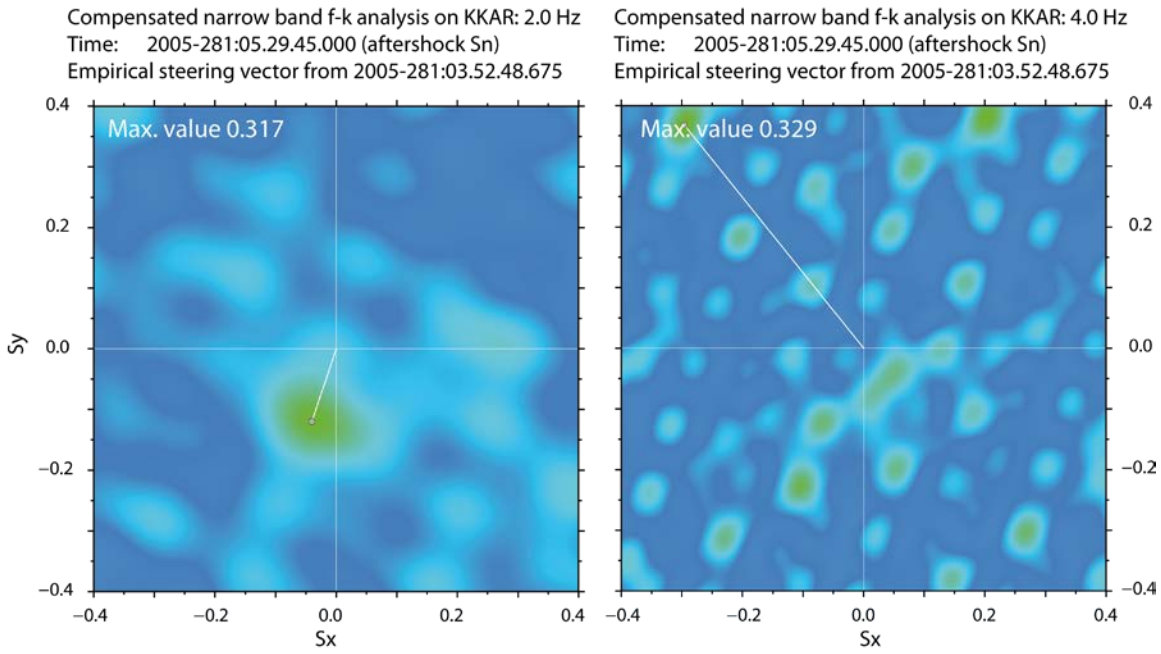

Fig. 4.3.4 Narrowband f-k grids for the Sn arrival of the aftershock generated after pre-steering by the empirical matched field steering vectors calculated from the main shock Pn arrival at 2 Hz (left) and 4 Hz (right).

Fig. 4.3.4 shows the corresponding relative f-k spectra calculated at the time of the Sn arrival at KKAR for the aftershock. The maximum value obtained is now significantly lower and, more significantly, is not obtained close to the zero slowness vector. This indicates that the approaching wavefront arrives from a direction other than that from which the empirical steering vector was calculated. For the 2 Hz signal, there is a single peak at approximately the vector difference between the Pn and Sn slowness vectors. At 4 Hz, there is no significant peak in the relative f-k spectrum and an essentially random point in slowness space is returned.

The conclusion from these figures is that we can probably devise an EMFP-based detection statistic to be applied as a primary detector. This matched field statistic will be high at times when the narrowband phase shifts across an array correspond well with those observed with the first arrival from the master event and low at times when the phase shifts correspond to a different direction of arrival. The following few figures illustrate the likely capabilities of such a primary detector.

The uppermost panel in Fig. 4.3.5 displays, for a 20 minute segment of this sequence, the relative power statistic as a function of time and frequency. We will refer to this as a pseudo-spectrogram. High values of the matched field statistic are shown in red and correspond clearly with the visible Pn-signal arrivals.

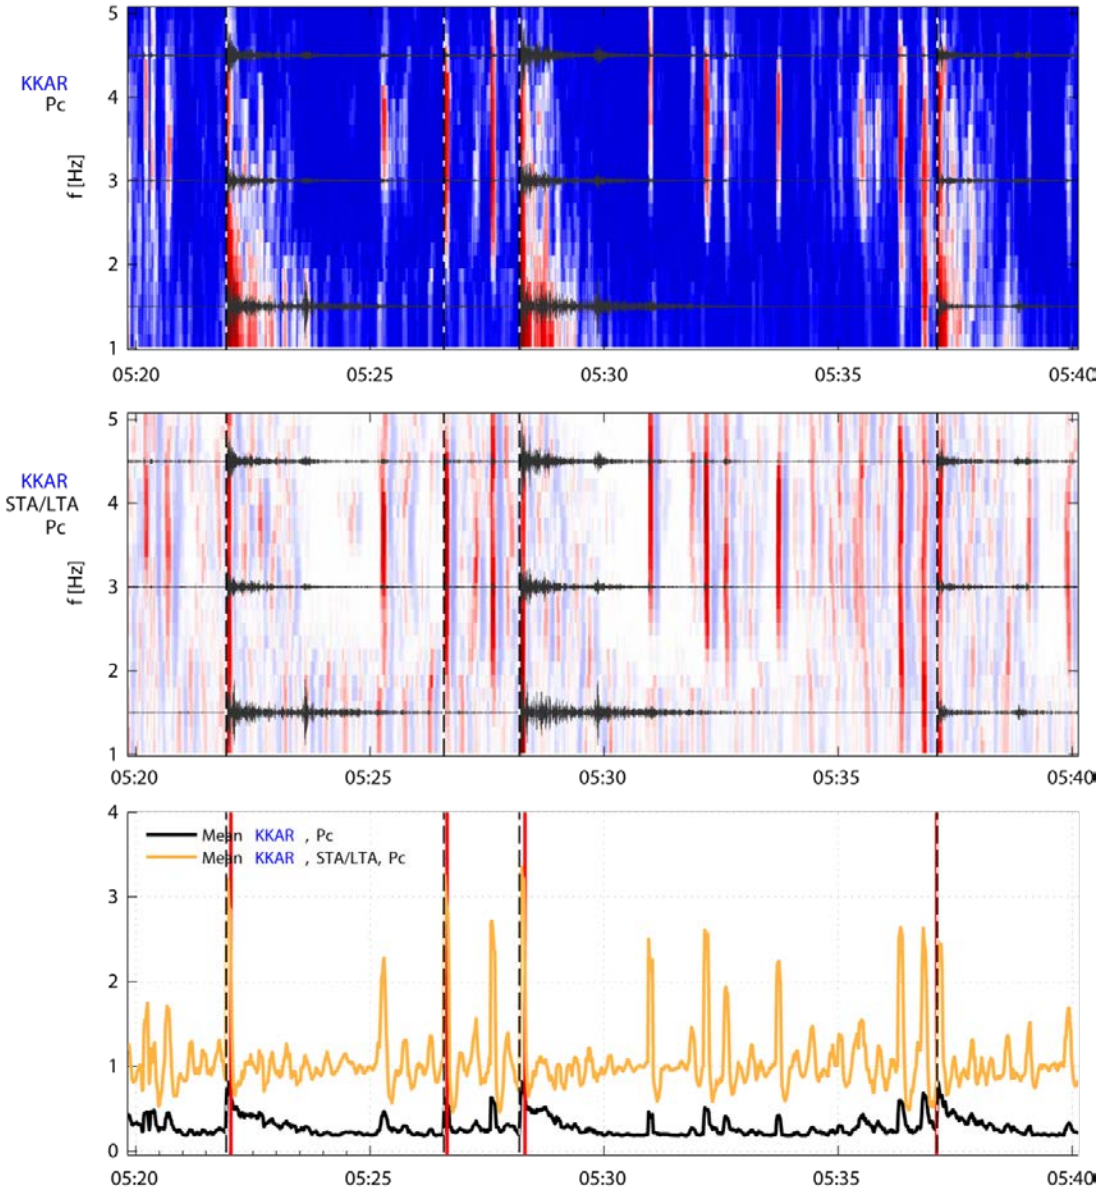

Fig. 4.3.5 The top panel pseudo-spectrogram shows the relative power matched field statistic Pc as a function of time and frequency. The waveforms centered on 1.5, 3, and 4.5 Hz display the KKAR Pn beam steered towards the source region, Butterworth band-pass filtered at the center frequencies indicated with bandwidth 1 Hz. The middle panel shows the STA/LTA ratio calculated for each frequency bin of the relative power pseudo-spectrogram displayed in the top panel. The length of the LTA segment is 5 seconds and it is separated from the 1 second long STA segment by a gap of 5 seconds. The bottom panel shows a black line for the relative power statistic Pr averaged over all frequencies of the top panel. Similarly, the orange line is frequency average of the STA/LTA ratio of the middle panel as a function of time. All times are UT on day 281 in 2005. In the top and middle panels, the vertical axis is the frequency in Hz. The vertical dashed lines correspond to arrival time estimates for events listed in the Reviewed Event Bulletin (REB) of the PTS.

The significant peaks in Sn-energy on the Pn-beams (that would generate a detection using an STA/LTA detection on the filtered waveforms themselves) do not correspond to high values of the matched field statistic. The use of the empirical matched field primary detector may therefore circumvent or reduce the need for f-k screening on triggers.

We note that for the larger signals, the vertical stripes which correspond to P arrivals cover the full (1-5 Hz) frequency band whereas the smaller signals only show matched field correspondence at the higher frequencies for which the signal-to-noise ratio (SNR) is better. At the lower frequencies, the good correspondence with the empirical matched field steering vectors continues well into the P-wave coda. At higher frequencies, the match with the empirical steering vectors is more transient and the relative power diminishes quite rapidly following the initial arrival. The STA/LTA-type transformation (per frequency band) displayed in the middle panel results in a function which peaks close to the signal onset and which diminishes into the signal coda far more rapidly than the matched field statistic itself. The transformed relative power gives the impression of vertical bars, covering a broad range of frequencies, at the times of the signal onsets. It is also clear from Fig. 4.3.5 that there is evidence of Pn phase arrivals from the source region of interest for many more events than are listed in the Reviewed Event Bulletin (REB) published by the Provisional Technical Secretariat (PTS) of the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) in Vienna. It should be noted though that KKAR is not a CTBTO seismic array and is significantly closer and significantly more sensitive for this particular sequence than any IMS arrays.

Fig. 4.3.6 displays the empirical matched field pseudo-spectrograms and the associated STA/LTA ratios for the IMS arrays BVAR, WRA, and ILAR, together with those for KKAR, and aligned with respect to the traveltimes to each of the arrays from the location of the main shock. Over the time-scale displayed, the vertical bars in each of the panels are almost perfectly aligned indicating a common source region for each of the events that generated these signals. Most of the vertical stripes in the KKAR panel are accompanied by similar stripes in the panel for BVAR, the Borovoye array in northern Kazakhstan. The stripes in the KKAR panel without corresponding signals at BVAR are signals that are limited to higher frequencies and therefore presumably of lower amplitude. The WRA array (Warramunga, Northern Territories, Australia) and the ILAR array (Eielson, Alaska, United States) share many of the same stripes displayed by the Kazakh arrays, although there are clearly signals at KKAR and BVAR that do not have a detection at the far more distant arrays.

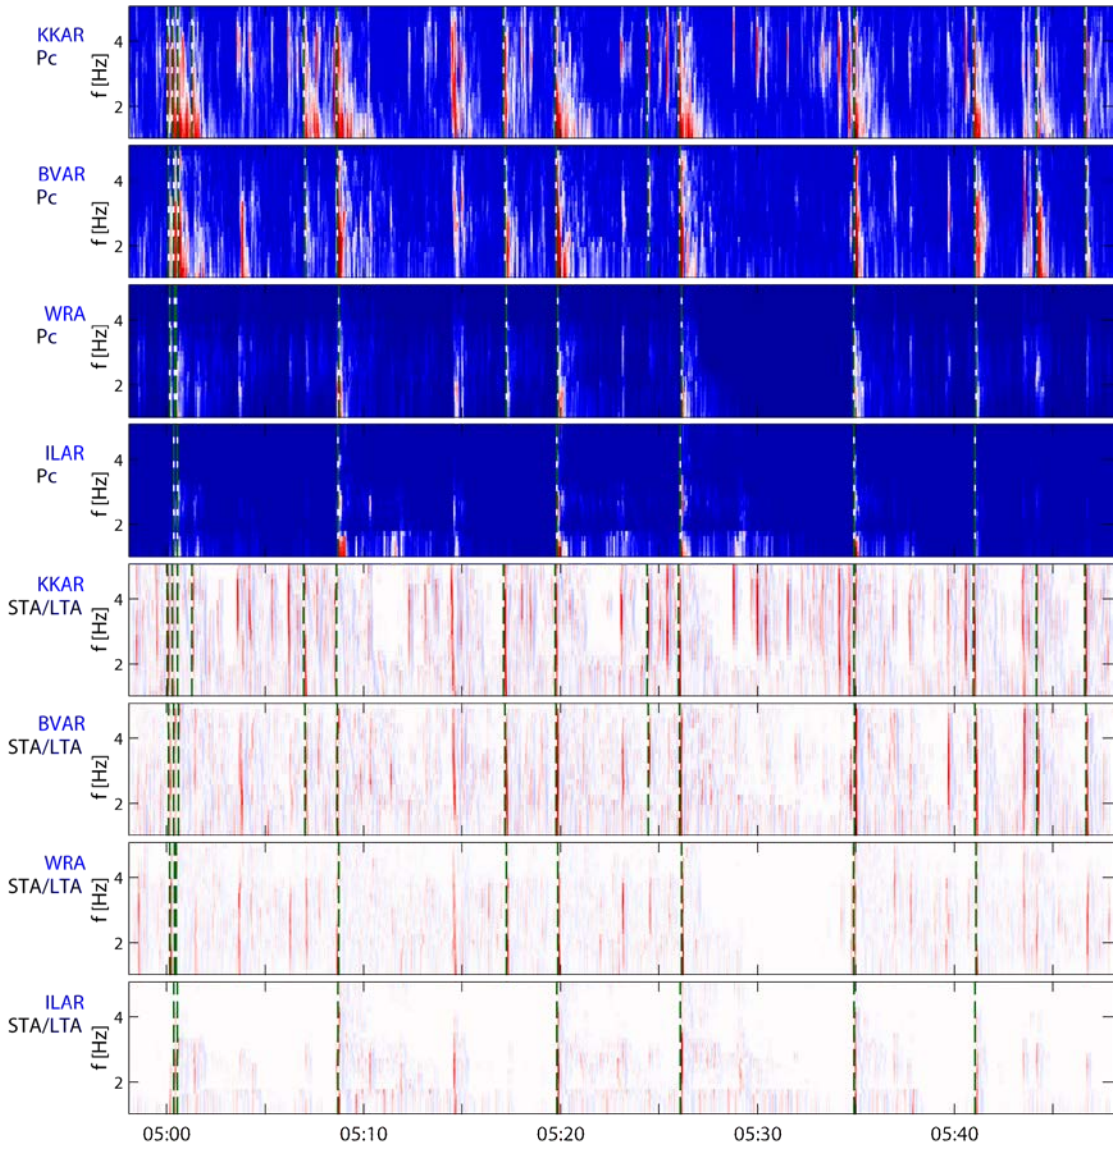

Fig. 4.3.6 The top four panels are pseudo-spectrograms for the KKAR, BVAR, WRA, and ILAR arrays for a 1 hour sequence on day 281 in 2005. For each array, empirical steering vectors were calculated for the first arrival from the main shock (this is a regional Pn phase for KKAR and teleseismic P for the other three arrays). All of the traces are back propagated to the origin time for the main event. The bottom four panels show the corresponding frequency-dependent STA/LTA ratios calculated in a similar manner as in the previous Fig. 4.3.5. The vertical dashed lines correspond to arrival time estimates for events listed in the Reviewed Event Bulletin (REB) of the PTS.

The time-frequency functions displayed in Fig. 4.3.6 can be transformed into scalar functions of time. A mean value across frequencies is the most straightforward option, although this would penalize smaller signals for which the local maxima do not span the

full range of frequencies. Fig. 4.3.7 displays such scalar functions for the four arrays displayed in Fig. 4.3.6 together with AKASG (Malin array, Ukraine), FINES (Lahti, Finland), and SONM (Songino, Mongolia).

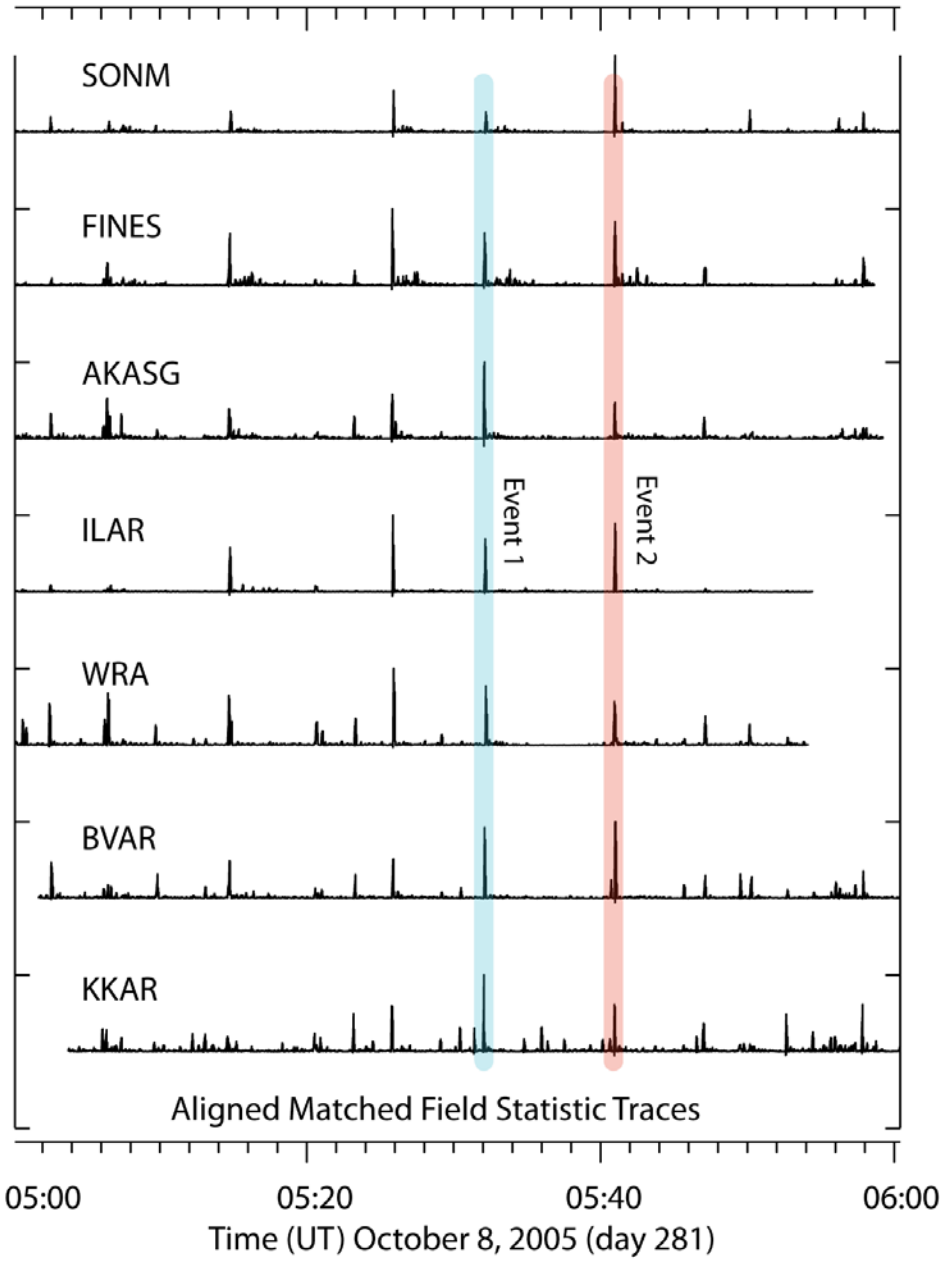

Fig. 4.3.7 Scalar functions derived from the STA/LTA transformations of the matched field pseudo-spectrograms for seven seismic arrays as labelled. Each of the traces is shifted backwards in time by the traveltime to the station from the main event (from which the EMFP template was defined). Two sets of peaks in these functions have been associated and assigned to events labelled "Event 1" and "Event 2".

While local maxima at the different arrays in Fig. 4.3.7 appear to be perfectly aligned over the time-scale displayed (one hour), measuring the times at which the maximum value is obtained reveals that the matched field scalar traces are in fact delayed by time-shifts that can be measured accurately enough to locate the events relative to each other Fig. 4.3.8.

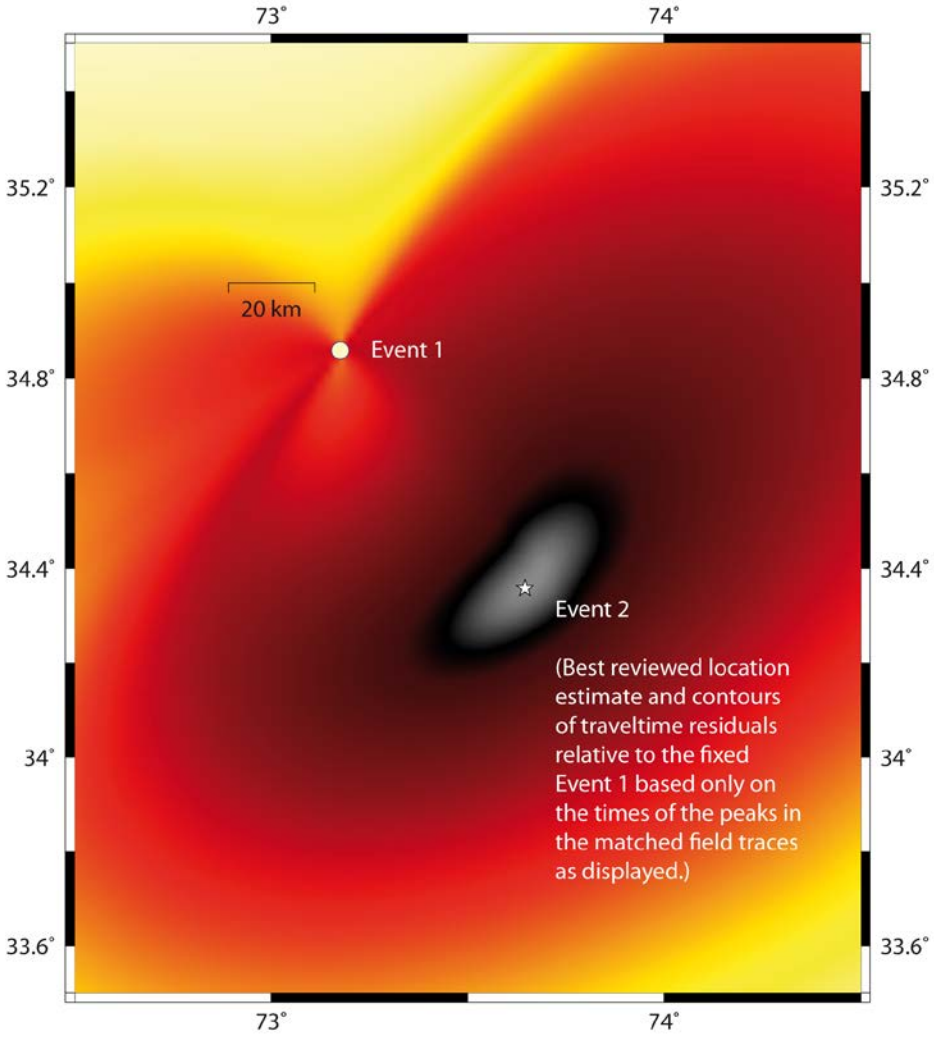

Fig. 4.3.8 If we fix the location of Event 1 to the point indicated by the white circle and measure the times of the corresponding maxima in the matched field scalars, we can calculate contours of misfit for the location of Event 2 by summing the residuals between predicted traveltime difference and measured traveltime difference (essentially a double-difference relative location estimate). Even using only the seven global seismic arrays displayed in the previous figure, we are able to constrain the location of Event 2 using only the times of the local maxima of the matched field output. The current best available analyst location estimate for Event 2 falls approximately at the center of the region with the lowest traveltime difference residual.

Finally, in Fig. 4.3.9, we demonstrate that EMFP provides not only a detector with a low false alarm rate, but also with a high sensitivity. Here we consider the KKAR array in Kazakhstan observing teleseismic P arrivals from the 23 October, 2011, Van earthquake in Eastern Turkey. A significant SNR is observed in the matched field scalar trace even when no significant SNR is seen in the waveforms themselves. (This is to say that we detect an arrival as defined by a recognized pattern in phase shifts between the different channels on the array with a far higher level of significance than we can detect an arrival based upon increased amplitudes of the signals themselves.) The authenticity of the arrivals at the distance KKAR array is confirmed by signals on a station far closer to the source region. Fig. 4.3.9 confirms in addition the superiority of the matched field detector over a simple correlator when a very large main shock is used as a master event.

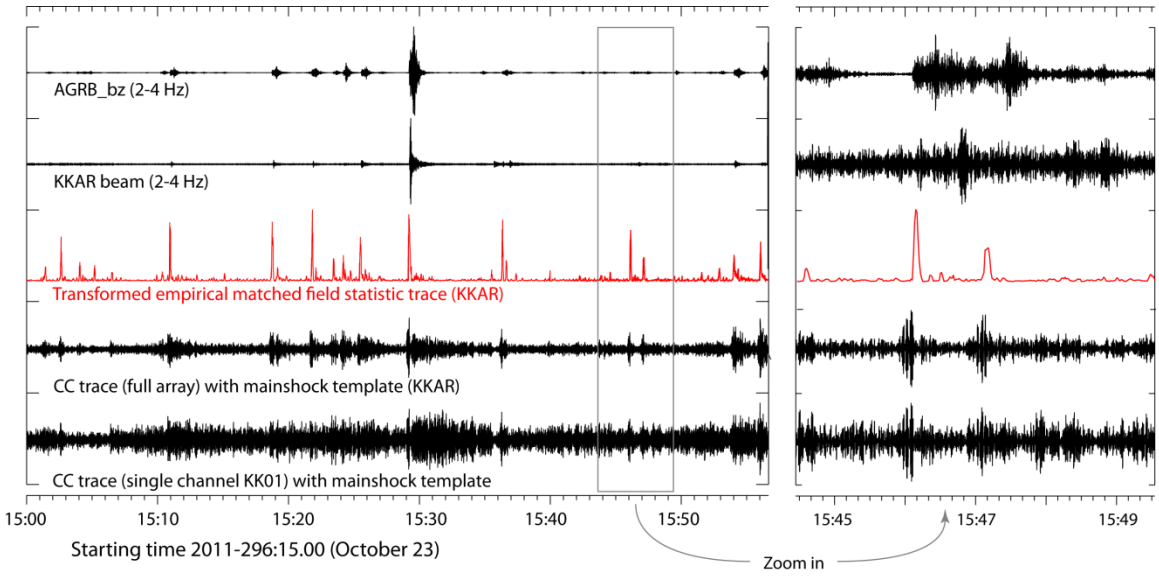

Fig. 4.3.9 Panel displaying a matched field scalar trace at KKAR (in red) together with the actual array beam (above) for the October 23, 2011, M=7.1 Van earthquake in Eastern Turkey. For the panel to the right, two clear detections are observed in the matched field trace for which no signal onsets are visible on the beam for the same array. In the very top trace is a seismogram from a local station (at a distance of approximately 1 degree), corrected for the traveltime from the source region, which displays a signal corresponding to the time of the matched field P-phase detection at KKAR. In the lowermost two traces are cross-correlation traces with the signal from the main shock: although a modulation is observed, the peaks are not well-defined in time or significantly above the background level.

4.4 Necessity of Screening Criteria for Template Definition

If the primary detector is a beam on a single array aimed towards the source region of interest, then the framework can be considered to be optimal for arrivals with a given slowness vector. However, signals from all directions are likely to be capable of generating triggers on this beam since a conventional beam is not capable of suppressing completely energy approaching from different backazimuths and angles of incidence. Many applications of the framework are likely to run over extended periods of time such that it is inevitable that, in the absence of a sufficiently robust screening algorithm, the number of correlation-type detectors will just escalate indefinitely spawning pattern detectors for signals with no relevance to the monitoring scenario.

This was found to be the case for the Storfjorden sequence (see Section 3.3) with the situation being exacerbated by the SPITS array's small aperture and consequent lack of signal suppression. The prolonged period of monitoring (several years) meant that vast numbers of signals from other directions resulted in triggers from the primary detector. The framework was modified to reject all primary triggers on arrays which did not fall within a predefined range of backazimuths and apparent velocities.

In the example of monitoring repeating mining explosions in Kazakhstan described in Section 3.4, we were only interested in setting up pattern detectors for the mines due north of the ABKAR array, and limits on the slowness of the primary detections were set as displayed in Fig. 4.4.1. Only the detection shown in the left hand panel would result in the generation of a pattern detector. Given the nature of the primary detectors, there was no way to accept primary detections for the Terensay mine without, for example, also accepting primary detections for the Koktau and Dombarovskiy mines (see Figure 3.4.1). As a result, pattern detectors resulted for all of these clusters although, as displayed in Fig. 3.4.2, the different sources were well-resolved by the framework into appropriate event clusters.

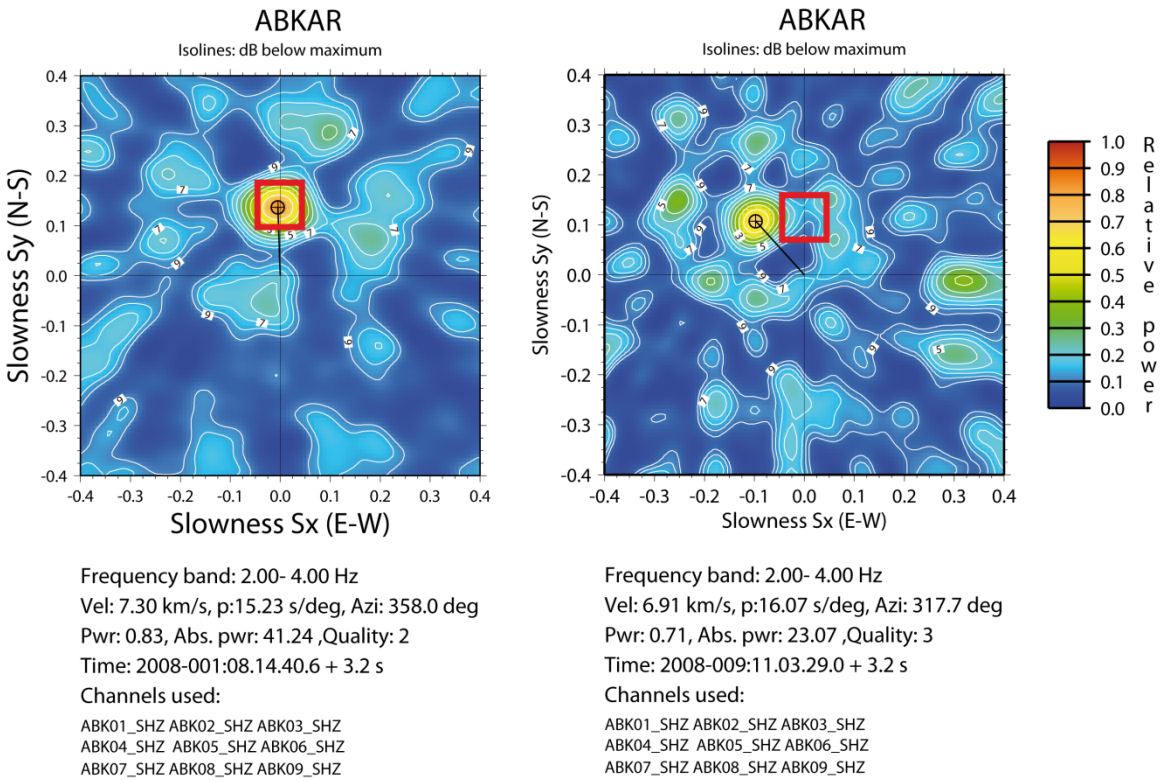

Fig. 4.4.1 f-k grids surrounding primary detector triggers on the ABKAR array for confirmed events from the Terensay mine (left) and the Chromtau mine (right). (Details of the locations of these mines relative to the ABKAR array are provided in Table 3.4.1 and Fig. 3.4.1) The red box denotes an approximate region of slowness space for which detections will be accepted for spawning pattern detectors for monitoring events at the Terensay mine. If f-k analysis indicates a direction not contained within this region of slowness space then the detection is discarded.

4.5 Mitigation of the effects of interfering signals in pattern detector templates

A common characteristic of rapidly evolving aftershock sequences is that events often occur so frequently that the time between two subsequent aftershocks can be significantly less than the signal duration at key stations. In correlation and subspace detectors, it is usually beneficial to take as long a waveform segment as possible so as to maximize the time-bandwidth product of the signal. Experience from the Kashmir sequence indicated that, unchecked, this property would result in large clusters consisting of detections on noise resulting from an interfering signal later in the waveform template.

An example is presented in Fig. 4.5.1 of how a fixed-length template which happens to be contaminated by a foreign signal at the end misses the signal of an event with which it shows great similarity. Moreover, a completely spurious detection occurs later on in the time segment. Fig. 4.5.2 indicates how, for the correlation detectors alone, this scenario is avoided by scaling each point of the waveform by the moving average of the absolute values: a kind of automatic gain control. Further details of this procedure are provided by Gibbons *et al*. (2012).

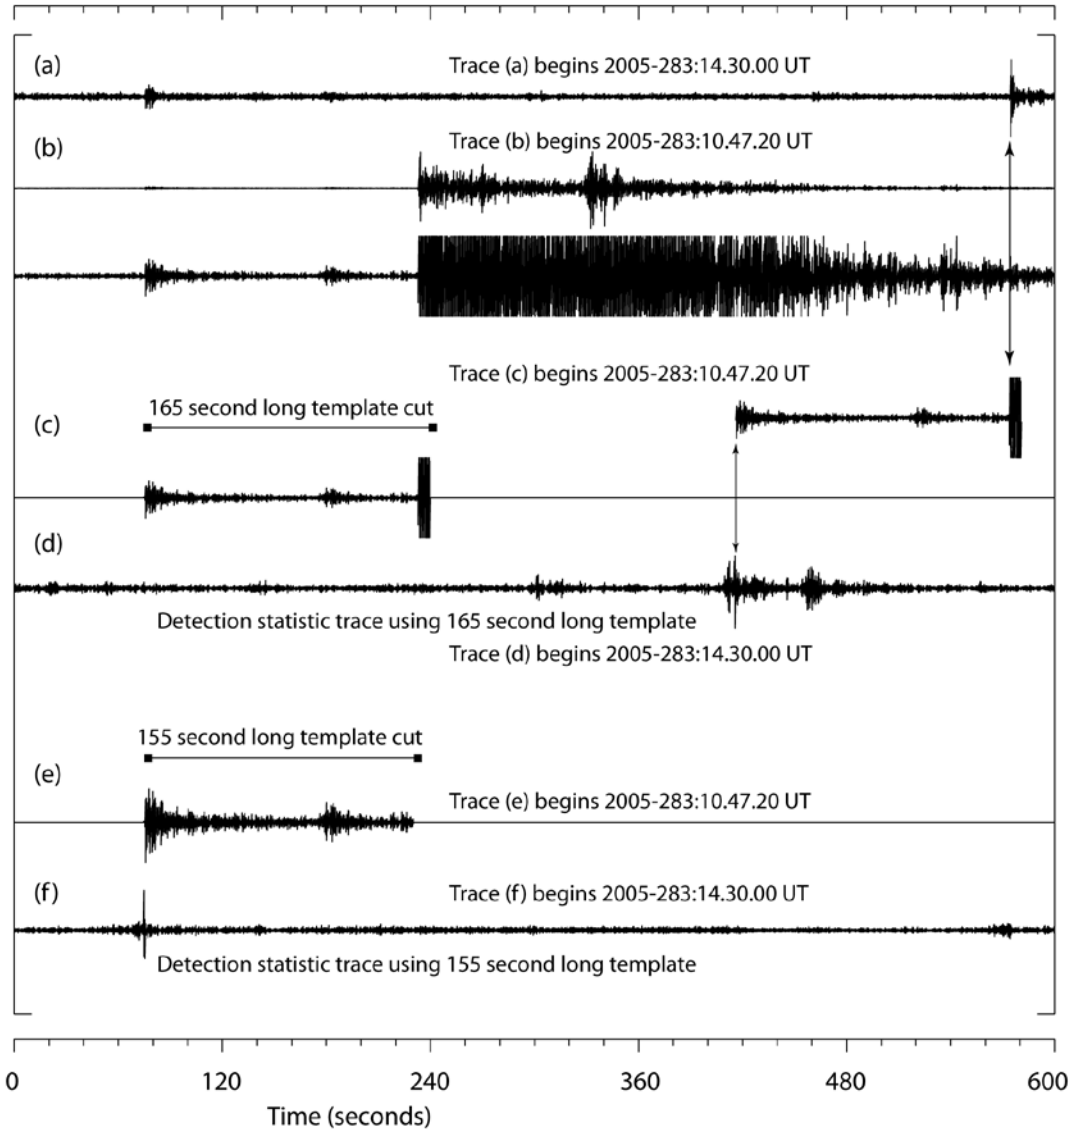

Fig. 4.5.1 The detection of spurious events and failure to detect a correlating event due to the contamination of the waveform template with a foreign signal within the signal template. A signal starting at a time 2005-283:10.48.35 (in trace b) correlates well with a later signal starting at a time 2005-283:14.31.15 (trace a). However, under the initial fixed-length template recipe, the master waveform also included a high amplitude signal beginning at a time 2005-283:10.51.13 (trace c). Correlating this template with the incoming waveforms results in the detection statistic (trace d) which completely fails to register the repeating event and produces a spurious detection later in the time-window. Selecting a 10 second shorter template (trace e) avoids the interfering signal and registers a significant correlation at the time of the repeating event signal.

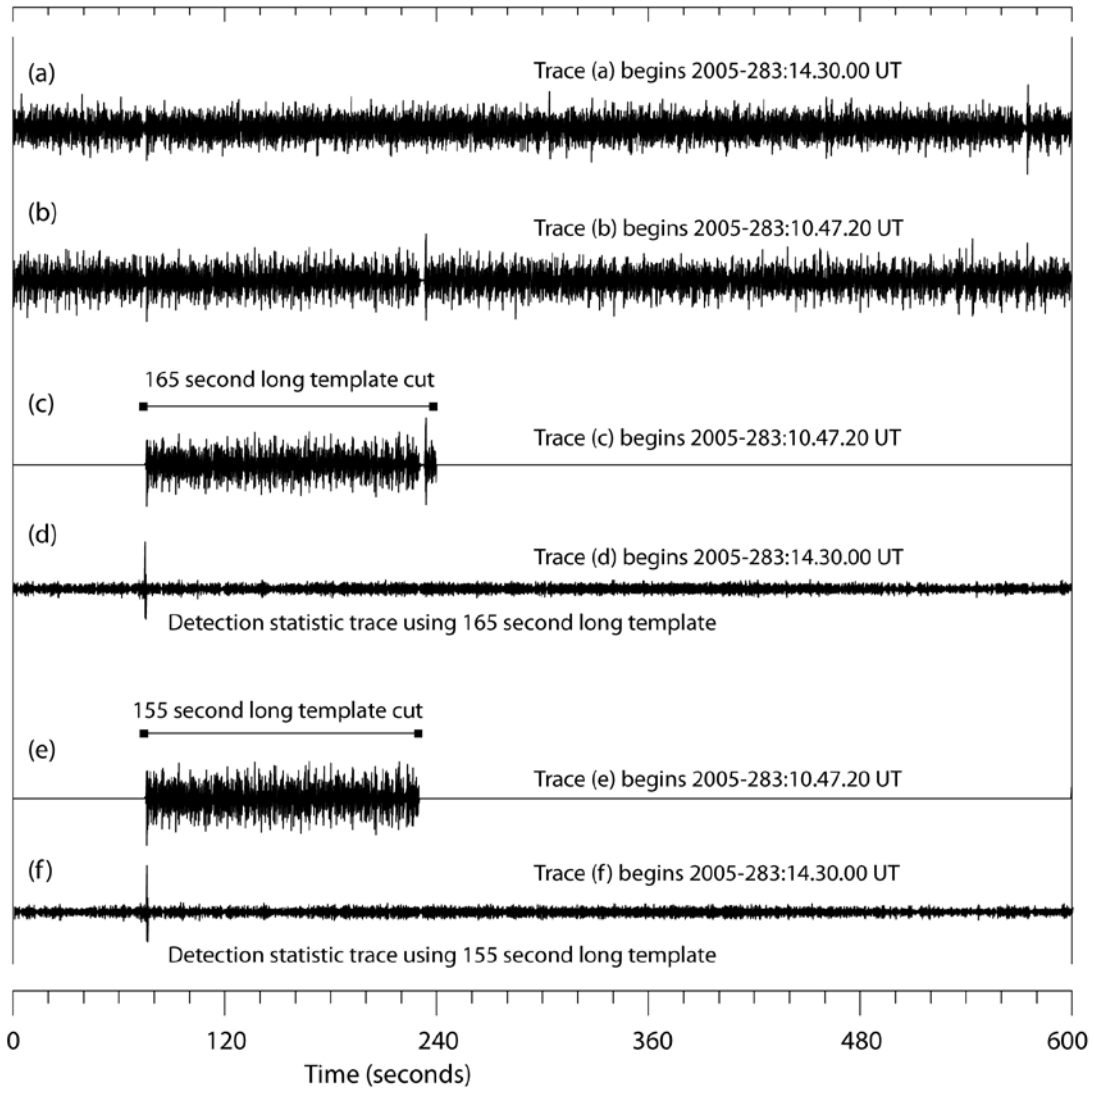

Fig. 4.5.2 An exact replication of the correlation calculations displayed in Figure 4.5.1 except that all filtered waveforms have been scaled by a moving average of the absolute values prior to the correlation. In this case, the influence of the interfering signal is marginal and the repeating event signal is detected by both templates. Also, no spurious detection is made later in the data segment.

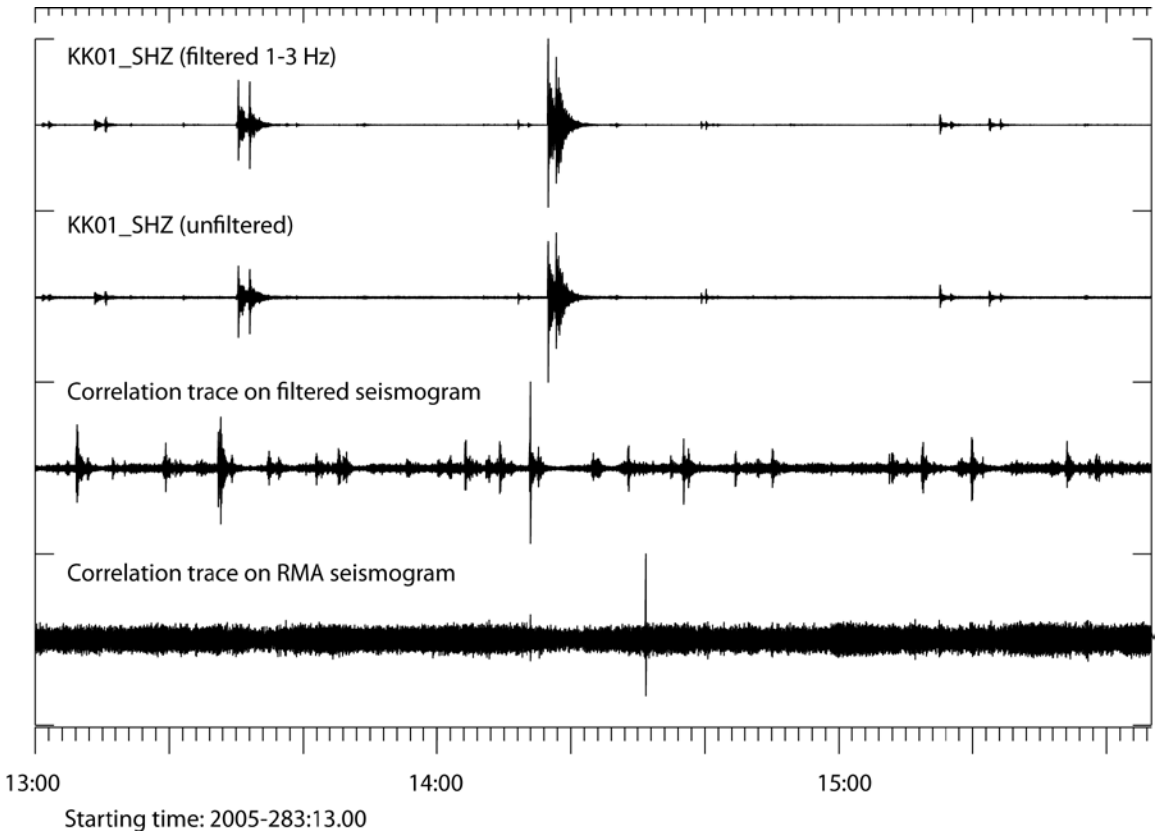

Fig. 4.5.3 The correlation traces over a long time interval for the 165 second long templates starting at time 2005-283:10.48.35 with and without a scaling by the moving average (in traces 1 and 2 from the bottom respectively). Only a single detection is made for the RMA trace whereas multiple detections are made when the correlation is performed on the unscaled traces.

It is illustrated in Fig. 4.5.3 how this pre-processing of the data significantly reduces the variability of the background level of the detection statistic and results in only a single detection at the time of the repeating signal. In comparison, numerous detections occur without the pre-processing and almost all are largely the result of the intrusive signal. Whereas the automatic gain control approach may be highly beneficial for the correlation detector, it would need to be handled in a completely different data stream in practice since the primary detectors would never trigger on these waveforms that displayed such little variation in amplitude. Instead the definition of the waveform templates are subject to a length which is allowed to vary between two pre-specified limits based upon the values of the time-center of the waveform $x(t)$, given by

$$t_0 = \frac{\int_{-\infty}^{\infty} t|x(t)|^2 dt}{\int_{-\infty}^{\infty} |x(t)|^2 dt}$$

and the time variance, given by

$$\sigma_t^2 = \frac{\int_{-\infty}^{\infty}(t-t_0)^2|x(t)|^2dt}{||x||_2^2}$$

In the framework, an iterative procedure is run whereby the template is shortened progressively until $t_0 + \sigma$ is less than the specified minimum template length. In the majority of cases considers this results in templates which are either excluding foreign signals altogether or which are not dominated by them. Examples are displayed in Figs. 4.5.4 and 4.5.5 of templates whereby the length is made to exclude new arrivals of energy significantly after the trigger time. The KKAR array is unusual for the Kashmir sequence in that the wavetrain contains both a Pn and an Sn arrival (at most observing arrays, only a P-wave is seen). In Fig. 4.5.5 it is clear that, with the parameters specified here, the Sn phase and coda is omitted from the template.

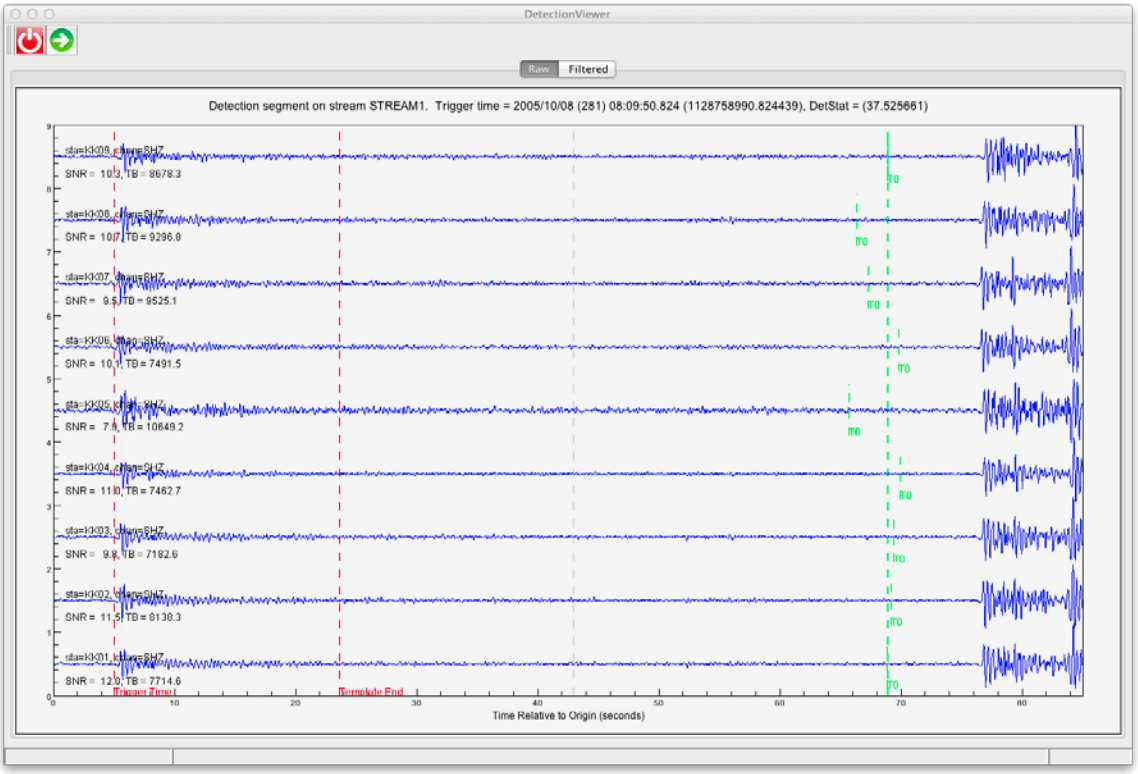

Fig. 4.5.4 Example 1 of setting a template according to the time center and time-variance.

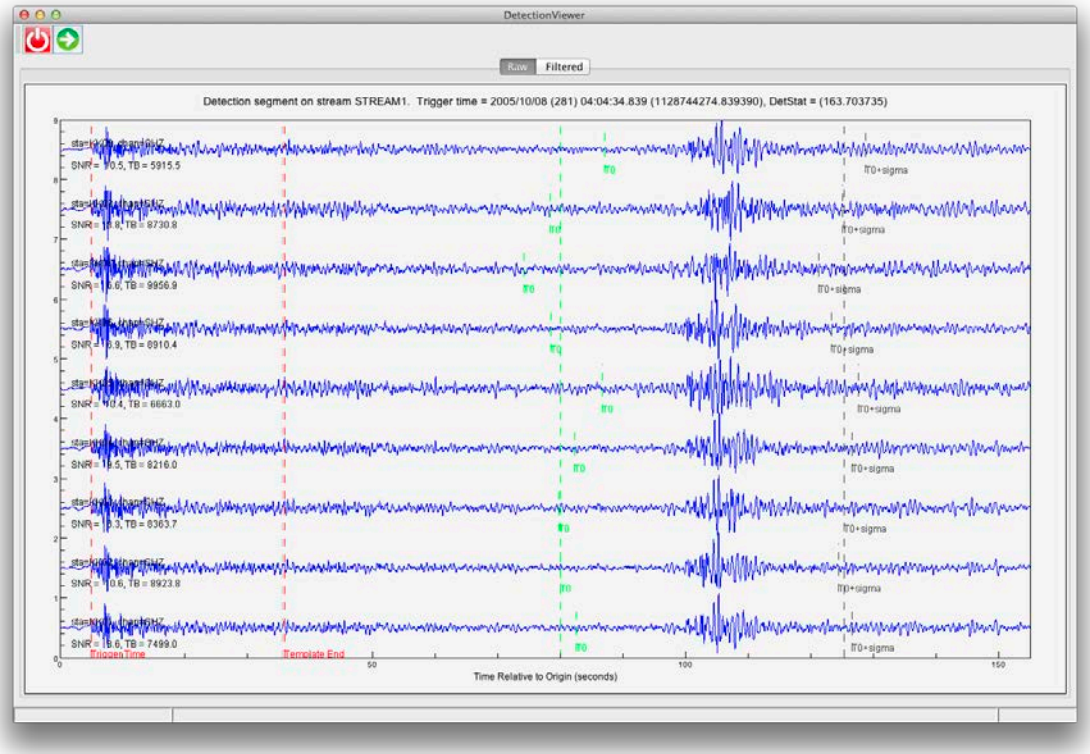

Fig. 4.5.5 Example 2 of setting a template according to the time center and time-variance.

5 CONCLUSIONS

A detection framework was proposed that would augment a routine processing pipeline with pattern detectors to identify and classify repeating waveforms in data streams from a network of seismic arrays. The framework has been implemented as proposed and tested by several institutes on a variety of different test cases. These range from extensive aftershock sequences from extremely large earthquakes recorded at large distances (Kashmir, 2005; Sumatra, 2010), moderate aftershock sequences observed at near-regional distances (Storfjorden, 2008), and mining blasts from multiple repeating sources at regional distances (Western Kazakhstan).

The performance of the framework, defined in terms of the proportion of seismicity which is classified correctly in clusters for rapid analyst interpretation, is far better for moderate sequences and seismicity recorded at regional distances than for aftershock sequences from exceptionally large earthquakes recorded at large distances. For the repeating mining blast scenario, large clusters of events were obtained and each was found to consist exclusively of signals generated by blasts at a single mine – confirmed by Ground Truth information. This essentially eliminates the need for an analyst to perform manual evaluation on single events for what is ultimately a screening process. For the Storfjorden sequence, the existing event bulletin (the NORSAR regional reviewed bulletin) was limited by analyst resources to events exceeding magnitude 2. The application of the detection framework reduced this threshold by more than a unit of magnitude over which a good overview of events, classified into clusters, is obtained. The framework provides a far more complete event catalog, of a quality which is of use in structural interpretation, at a modestly increased workload.

The Kashmir and Sumatra cases are more challenging given the exceptionally large source regions covered by the aftershocks. The spatial separations between the hypocenters of the largest events (i.e. those that are well recorded at teleseismic distances) are simply too large to result in significant waveform semblance, in the frequency bands of interest, between the signals from subsequent events. In these cases, the framework often identifies clusters of events in pockets of seismicity within the extended source region: often at lower magnitudes than appeared in the existing catalogs. The formation of high-rank subspace detectors can sweep up large proportions of the seismicity although not necessarily in a way that reduces significantly the burden of analyst interpretation. We suggest the use of a subspace-measure of waveform similarity that may perform better than the classical correlation coefficient for forming and evaluating clusters in such cases.

We have demonstrated a need to mitigate the problems caused by overlapping events which, if not corrected, result in both missed detections and spurious detections. The framework has mechanisms for both automatic and manual modification of waveform templates to avoid or limit the influence of interfering signals. We have, in addition, submerged a signal of monitoring interest into an array data-stream during an extensive aftershock sequence and demonstrated that this signal is not screened out by the pattern detector components of the framework.

We propose empirical matched field processing (EMFP) on single array streams as a sensitive primary detector for generating triggers at, and only at, the times of arrivals from events in the region of interest. We have demonstrated that EMFP, which recognizes the pattern of narrowband phase shifts over an array for a previously observed incident phase, can display a significant SNR even when the SNR on the corresponding classical beam is low. The matched field detector shows promise for generating output that may be processed incoherently over multiple arrays for the creation of robust event hypotheses over an extended source region.

REFERENCES

Ali, Z., Qaisar, M., Mahmood, T., Shah, M. A., Iqbal, T., Serva, L., Michetti, A., and Burton, P. W. (2009) "The Muzaffarabad, Pakistan, earthquake of 8 October 2005: surface faulting, environmental effects and macroseismic intensity" Geological Society, London, Special Publications, Vol. 316, No. 1. (1 January 2009), pp. 155-172, doi:10.1144/sp316.9

Bendick, R., Bilham, R., Khan, M. A, Khan, S. F. (2006) "Slip on an active wedge thrust from geodetic observations of the 8 October 2005 Kashmir earthquake" Geology, Vol. 35, No. 3. (01 March 2006), pp. 267-270, doi:10.1130/g23158a.1

Gibbons, S. J., Kværna, T., and Ringdal, F. (2010) "Considerations in Phase Estimation and Event Location Using Small-aperture Regional Seismic Arrays" Pure and Applied Geophysics, Vol. 167, No. 4. (1 May 2010), pp. 381-399, doi:10.1007/s00024-009-0024-1

Gibbons, S. J., Ringdal, F., and Kværna, T. (2012) "Ratio-to-moving-average seismograms: a strategy for improving correlation detector performance" Geophysical Journal International, Vol. 190, No. 1. (01 July 2012), pp. 511-521, doi:10.1111/j.1365-246x.2012.05492.x

Golub, G. H. and C. F. Van Loan (1996), Matrix Computations, Johns Hopkins University Press, ISBN-13: 978-0801854149

Harris, D. (1991) "A waveform correlation method for identifying quarry explosions" Bulletin of the Seismological Society of America, Vol. 81, No. 6. (1 December 1991), pp. 2395-2418

Harris, D. B. and Dodge, D. A. (2011) "An Autonomous System for Grouping Events in a Developing Aftershock Sequence" Bulletin of the Seismological Society of America, Vol. 101, No. 2. (1 April 2011), pp. 763-774, doi:10.1785/0120100103

Harris, D. B. and Kværna, T. (2010) "Superresolution with seismic arrays using empirical matched field processing" Geophysical Journal International, Vol. 182, No. 3. (September 2010), pp. 1455-1477, doi:10.1111/j.1365-246x.2010.04684.x

Jayangondaperumal, R., and Thakur, V. C. (2008) "Co-seismic secondary surface fractures on southeastward extension of the rupture zone of the 2005 Kashmir earthquake" Tectonophysics, Vol. 446, No. 1-4. (01 January 2008), pp. 61-76, doi:10.1016/j.tecto.2007.10.006

Jouanne, F., Awan, A., Madji, A. , Pêcher, A. , Latif, M., Kausar, A. , Mugnier, J. L., Khan, I., and Khan, N. A. (2011) "Postseismic deformation in Pakistan after the 8 October 2005 earthquake: Evidence of afterslip along a flat north of the Balakot-Bagh thrust" Journal of Geophysical Research, Vol. 116, B07401, doi:10.1029/2010JB007903

Junek, W., Vandemark, T., Saults, T., Harris, D., Dodge, D., Ichinose, G., Poffenberger, A., and Kemerait, R. (2013) "Integration of Empirical Signal Detectors into the Detection and Feature Extraction Application at the United States Data Center" Poster T3-P63 at Science and Technology 2013 conference, Vienna, 17-21 June 2013

Junek, W., N., Kværna, T., Pirli, M., Schweitzer, J., Harris, D. B., Dodge, D. A., Woods, M. T. (2014, in preparation) "Inferring Aftershock Sequence Properties and Tectonic Structure Using Empirical Signal Detectors"

Kaneda, H., Nakata, T., Tsutsumi, H., Kondo, H., Sugito, N., Awata, Y., Akhtar, S. S., Majid, A., Khattak, W., Awan, A. A., Yeats, R. S., Hussain, A., Ashraf, M., Wesnousky, S. G., and Kausar, A. B. (2008) "Surface Rupture of the 2005 Kashmir, Pakistan, Earthquake and Its Active Tectonic Implications" Bulletin of the Seismological Society of America, Vol. 98, No. 2. (01 April 2008), pp. 521-557, doi:10.1785/0120070073

Kennett, B.L.N., Engdahl, E.R., & Buland R. (1995) "Constraints on seismic velocities in the Earth from travel times" Geophys. J. Int., vol. 122, 108-124

Mandal, P., Chadha, R. K. , Kumar, N. , Raju, I. P. , and Satyamurty, C. (2007) "Source Parameters of the Deadly Mw 7.6 Kashmir Earthquake of 8 October, 2005" Pure and Applied Geophysics, Vol. 164, No. 10. (4 October 2007), pp. 1963-1983, doi:10.1007/s00024-007-0258-8

Murphy, J. R. , Rodi, W. , Johnson, M. , Sultanov, D. D. , Bennett, T. J. , Toksöz, M. N. , Ovtchinnikov, V., Barker, B. W., Reiter, D. T. , Rosca, A. C. , Shchukin Y. (2005) "Calibration of International Monitoring System (IMS) Stations in Central and Eastern Asia for Improved Seismic Event Location" Bulletin of the Seismological Society of America, Vol. 95, No. 4. (01 August 2005), pp. 1535-1560, doi:10.1785/0120040087

Ogata, Y. (1988) "Statistical models of earthquake processes and residual analysis for point processes" Journal of the American Statistical Association, Vol. 83, No. 401, pp- 9-27

Parsons, T., Yeats, R. S., Yagi, Y., Hussain, A. (2006) "Static stress change from the 8 October, 2005 M = 7.6 Kashmir earthquake" Geophysical Research Letters, Vol. 33, No. 6. (21 March 2006), L06304, doi:10.1029/2005gl025429

Pathier, E. , Fielding, E. J. , Wright, T. J. , Walker, R. , Parsons, B. E. , and Hensley, S. (2006) "Displacement field and slip distribution of the 2005 Kashmir earthquake from SAR imagery" Geophysical Research Letters, Vol. 33, No. 20. (24 October 2006), L20310, doi:10.1029/2006gl027193

Pirli, M., J. Schweitzer, and B. Paulsen (2013), The Storfjorden, Svalbard, 2008-2012 aftershock sequence: Seismotectonics in a Polar environment, Tectonophysics, doi:10.1016/j.tecto.2013.05.010

Regional Seismic Travel Times (RSTT) Software and documentation available from http://sourceforge.net/projects/rstt-software/

Slinkard, M. E., Carr, D. B., Young, C. J. (2013) "Applying Waveform Correlation to Three Aftershock Sequences" Bulletin of the Seismological Society of America, Vol. 103, No. 2A. (01 April 2013), pp. 675-693, doi:10.1785/0120120058

Soltanolkotabi, M. E. Elhamifar and E. J. Candes (2013), Robust SubspaceClustering, arxiv.org, arXiv:1301.2063

Tahir, M., and Grasso, J. R. (2014) "Aftershock Patterns of Ms>7 Earthquakes in the India–Asia Collision Belt: Anomalous Results from the Muzaffarabad Earthquake Sequence, Kashmir, 2005" Bulletin of the Seismological Society of America, Vol. 104, No. 1., pp. 1-23, doi:10.1785/0120120158

Yan, Y., Pinel, V., Trouvé, E., Pathier, E., Perrin, J., Bascou, P., Jouanne, F. (2013) "Coseismic displacement field and slip distribution of the 2005 Kashmir earthquake from SAR amplitude image correlation and differential interferometry" Geophysical Journal International, Vol. 193, No. 1. (01 April 2013), pp. 29-46, doi:10.1093/gji/ggs102

APPENDIX - RELOCATION OF EVENTS IN THE KASHMIR AFTERSHOCK SEQUENCE

The M=7.6 earthquake on October 8, 2005, close to the city of Muzaffarabad in Pakistan administered Kashmir (Mandal *et al*., 2007) was one of the deadliest earthquakes of all time with approximately 100,000 fatalities. The following aftershock sequence was exceptionally long and intense in comparison with other quakes in the region with similar magnitudes (Tahir and Grasso, 2014). The locations of the aftershocks are of crucial importance in understanding the structure and tectonics of the region (e.g. Ali *et al*., 2009; Bendick *et al*., 2006; Jayangondaperumal and Thakur, 2008; Jouanne *et al*., 2011; Kaneda *et al*., 2008; Parsons *et al*., 2006; Pathier *et al*., 2006; Yan *et al*., 2013) and in the evaluation of algorithms for the detection and classification of seismic waveforms (e.g. Slinkard *et al*., 2013, in addition to the current study).

The available seismic bulletins for this particular aftershock sequence are not of especially high quality for a number of reasons which will be discussed in this brief summary. Two of the most complete event bulletins for the sequence are displayed in Fig. A.1. (Note that the REB solutions are also listed in the bulletin of the ISC and are therefore available to the general seismological community, regardless of access to IDC products.)

Considering first the REB (Fig. A.1, right), event locations are constrained only by phase readings from IMS stations. Due to political reasons, there are/were no IMS stations in India, Pakistan, or China, and key stations to the west and southwest (the GEYT array in Turkmenistan and the TORD array in Niger) were not yet in operation. The absence of array stations in sub-Saharan Africa means that a significant range of backazimuths are missing for all but the very strongest events. To the north, the closest IMS station which was certified at the time of the sequence is the Makanchi array (MKAR) in Kazakhstan although, at its far regional distance, this array is one of the poorest for recording events in this sequence (the arrivals are emergent and the coda long meaning that many arrivals are hidden by the coda of previous events). The AAK station in Kyrgyzstan was not yet in IDC operations (certified in 2007) and the very sensitive AKTO station in Kazakhstan was certified in November 2005 hence missing the main shock and the most significant aftershocks. The ZALV seismic array in Russia was also not completed (certified December 2006) although the 3-component station ZAL which it replaced was particularly efficient at recording P-arrivals for this sequence. The quiet ASAR array in Australia was unfortunately not in operations at the time although WRA (Warramunga) was and is one of the best array stations globally for monitoring this region. The station

coverage in northern Europe is excellent and very efficient propagation of high frequency energy to the stations in Scandinavia was beneficial for event forming at lower magnitudes. BVAR, the Borovoye array in Kazakhstan, is also very sensitive for this sequence but, being an auxiliary array, could not be used in the IDC's detection process (we of course are not subject to such restrictions). The MJAR array in Japan covers an otherwise sparse region of backazimuths and, despite its relatively high level of background noise and complicated site effects, is a relatively sensitive station for this sequence. The ILAR and YKA stations in Alaska and northern Canada respectively are important contributors to the REB for this sequence. The depths of almost all events in the REB are fixed to zero.

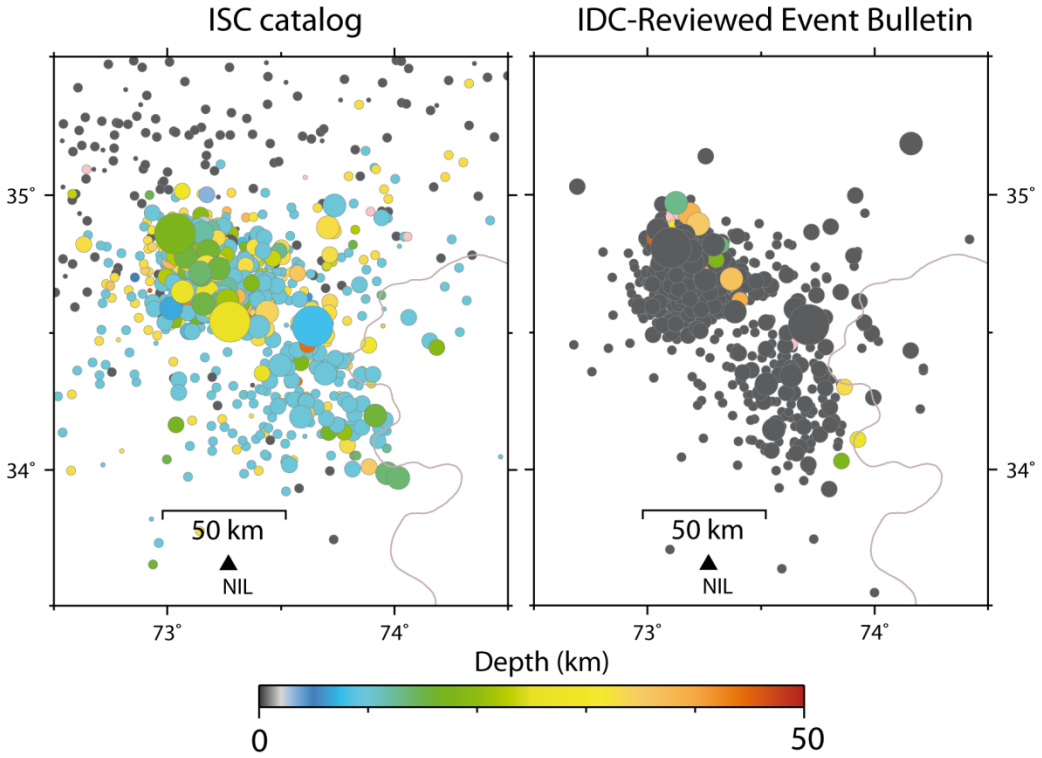

Fig. A.1 *Events from the bulletins of the International Seismological Center (www.isc.ac.uk, left) and the Reviewed Event Bulletin (REB) of the International Data Center (IDC) for the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) in Vienna (right) between October 8, 2005, and December 31, 2005.*

The ISC catalog is dominated, especially towards the start of the sequence, by event locations and phase readings from the REB. Many apparently quite large events are constrained only by teleseismic P-arrivals at, for example 4 or 5 teleseismic IMS arrays: the signals at other key stations being hidden in the noisy coda of earlier events. Some, although not all, of the ISC solutions are augmented by readings from non-IMS stations

picked by other agencies, and the vast majority of events have an ISC re-estimate of the location. The depths of these events are rarely fixed although, in the absence of strong seismological constraints, vary significantly with no patterns obvious in the left panel of Fig. A.1. A number of black dots are observed with an apparently random scatter among the ISC solutions; these are almost all single-array location estimates obtained from the KKAR array in Kazakhstan (using only the Sn and Pn phases at this array together with the backazimuth estimates obtained). The backazimuth estimates for the S-phases in particular are extremely prone to error on this vertical component seismic array and small differences in backazimuth estimates at 9 degrees distance can make significant differences to the resulting event locations. Significantly, the Nilore station (NIL) in Pakistan, by far the closest digital station whose data is readily available to the seismological community, does not appear anywhere in the ISC bulletin. While the NIL waveforms are severely clipped for the largest events, the first P-arrivals can be read well for most events and, several hours into the sequence, both S and P phases can be read with quite a high level of accuracy. Almost all of the events in the ISC catalog formed only from single array measurements at KKAR are seen clearly in the data at Nilore and combining arrival times at NIL, together with other available phase readings, should place important additional constraints on these events.

The stations of KNET (including AAK) record highly impulsive Pn arrivals for most events and some stations, notably AML, also record very clear Sn arrivals. Fortuitously, the MANAS profile (IRIS network XP: see Fig. 3.1.1) was deployed at the time and several of these stations recorded the vast majority of events very well with impulsive onsets for both Pn and Sn arrivals. While the azimuthal gap is not closed greatly by these stations, the constraints imposed upon epicentral distance with the enormous number of additional arrivals are significant.

It was deemed necessary to perform a complete relocation of the events in the sequence in order to be able to assess the performance of the detection framework on this sequence. Unfortunately, the additional high quality regional phases cannot simply be added to the inversion uncritically. As discussed by, for example, Murphy *et al*. (2005), station calibrations are required for the addition of regional phases to the location procedure. The slow, deep, and heterogeneous crust in the region results in significant event mislocation if the erroneous regional traveltimes are used uncorrected (see Fig. A.2). The RSTT (Regional Seismic Travel Times) software is based upon a 3-D model of crustal velocities and, in principle, we could simply use RSTT traveltimes for stations out to 15 degrees and follow classical location procedures. However, given the smaller but non-negligible bias in many of the teleseismic traveltimes, together with the unevaluated applicability of RSTT curves to the regional paths in question, it was

decided instead to attempt a multiple event, iterative, procedure to solve for time correction terms for all of the phases present, both regional and teleseismic.

This iterative procedure started by taking a number of the largest events with many teleseismic P-arrivals with excellent azimuthal coverage. For each event, a grid search was carried out, for a fixed depth of 20 km, covering the region displayed in the grids in Fig. A.2. For a given event, at each point in the grid, rays were raced out using the ak135 velocity model (Kennett *et al*., 1995) for all of the stations for which teleseismic P arrival times had been associated and an origin time was determined for that location which generated the lowest 1-norm residual vector. For the grid point at which the lowest 1-norm residual was found, the observed minus predicted traveltime differences were calculated for each phase. After all events had been processed, the residuals for each phase were examined and, if a significant and robust bias was identified, an empirical time correction for that phase was defined as the median traveltime residual. Using this set of empirical traveltime corrections, a new iteration was carried out with a new grid search being performed for each event. This would typically result in a slightly perturbed location estimate and a lower 1-norm traveltime residual. A new set of traveltime residuals were calculated and a second iteration of empirical traveltime corrections determined, always using the median of the distribution. Note that the same correction was assumed for the entire region displayed in Fig. A.2: no provision for variability as a function of source location was allowed. After a satisfactory convergence of the empirical teleseismic P traveltime corrections, using the best fixed depth location estimates, correction terms for the regional phases were calculated based upon the median values of the traveltime residuals. No iteration was carried out for the regional phase empirical traveltime corrections; the locations of these defining events were fixed based upon the teleseismic phases alone.

Fig. A.2 shows the 1-norm traveltime residuals for one particular event which was recorded with exceptionally many phase picks, both regional and teleseismic. The left panel indicates what these regional phase picks would do to the location estimate if not corrected: most of the regional stations are to the north and the traveltimes modelled are too short (i.e. the crust in this region is far slower than the ak135 model suggests). The longer times taken by these regional phases mean that, without applying the corrections, the locations are pushed too far south. Applying the traveltime calibrations brings the event back to the approximate location obtained using teleseismic phases only and has a 1-norm residual of the same order.

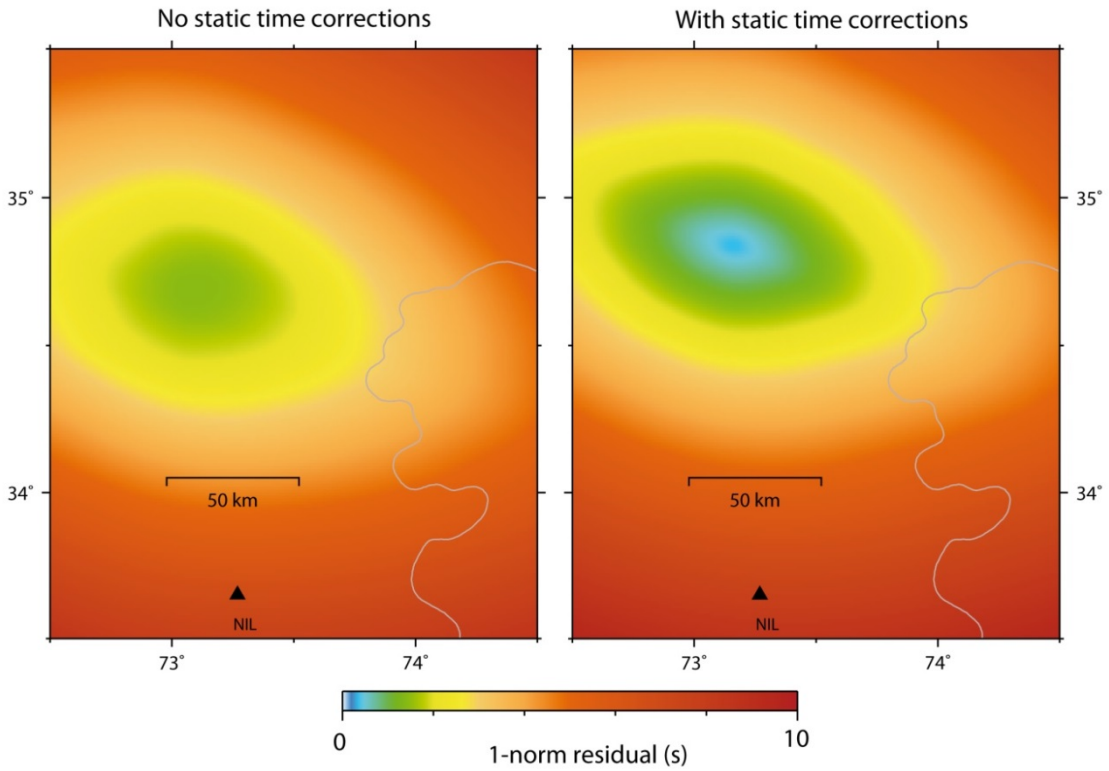

Fig. A.2 *Contours of 1-norm time-residuals for arrivals (both at regional and teleseismic distances) for a single event in the sequence (Event 1 in Figure 4.3.8) both with and without empirical time corrections solved for in this study. Using only teleseismic P-phases results in a location estimate close to that displayed in the right panel, but the addition of many uncorrected regional phases at stations with an unfavorable geometrical distribution relative to the source region pushes the location almost 20 km to the southwest and increases significantly the traveltime residual and hence the confidence in the event location estimate. Time corrections have been calculated for all stations but are far larger for the regional phases than for the teleseismic phases. When the time calibrations are imposed, the time-residual decreases greatly. The same time corrections are applied to all events in the sequence and are likely to be especially significant for smaller events for which we only have recordings at regional distances.*

The left panel of Fig. A.3 shows the relocations for all of the REB events using the grid search method, applying the traveltime calibrations calculated. The spread of the aftershock sequence has decreased considerably from the distributions displayed in Fig. A.1. There are a number of assumptions made that may make the calibrations invalid, for example that the values of the time correction terms estimated for a fixed depth of 20 km are valid for all depths, and that the calibrations do not change over the region displayed. (This is quite likely to be almost true for the teleseismic phases, and less likely to be a valid assumption for, for example regional phases at station NIL.) However, the calibrations obtained appear to be internally consistent to the extent that events which

are associated with both regional and teleseismic phases can now be located using either only regional phases or only teleseismic phases and produce consistent location estimates.

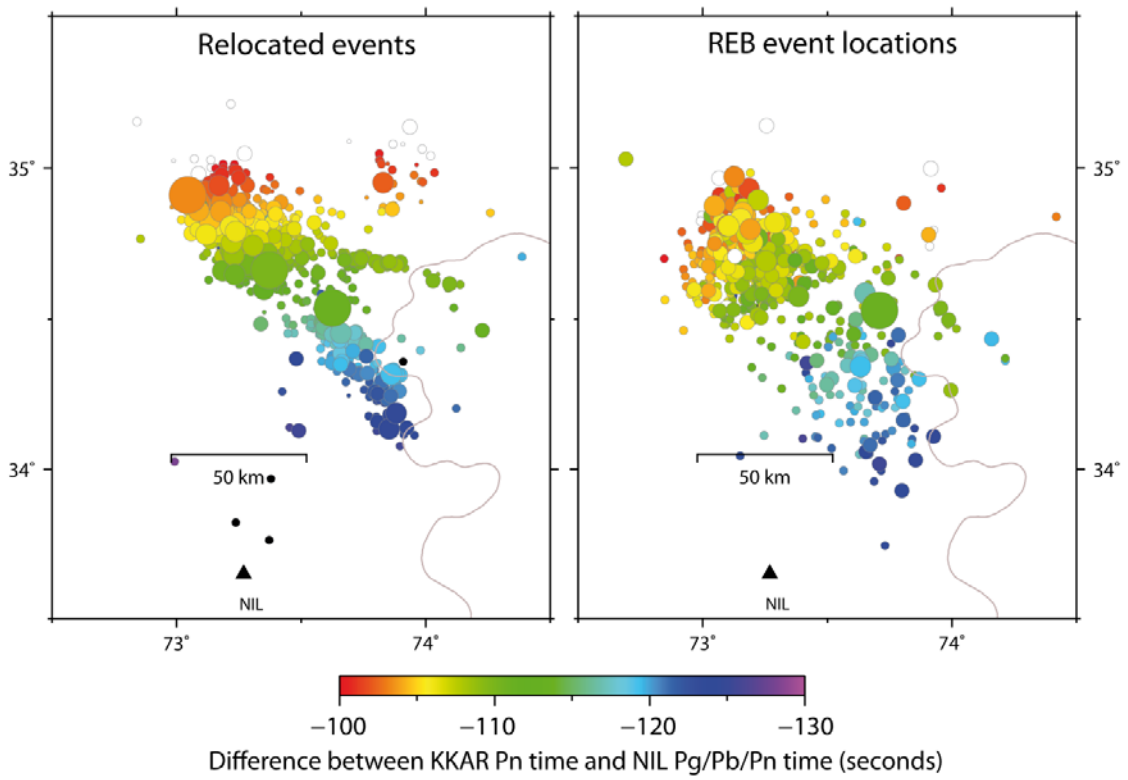

Fig. A.3 *A comparison between events relocated in the current study (left) and locations from the REB (right). The colors indicate the delay in seconds between the first P-arrival at NIL to the south and KKAR to the north for both distributions of event location estimates. While this does not provide an independent quality check on the event location estimates (since both stations were used in the event relocation), there is clearly a far more consistent pattern for the relocated events.*

SYMBOLS, ABBREVIATIONS, AND ACRONYMS

AFTAC – Air Force Technical Applications Center
CTBT – Comprehensive Nuclear-Test-Ban Treaty
CTBTO – Comprehensive Nuclear-Test-Ban Treaty Organization
EMFP – Empirical Matched Field Processing
IDC – International Data Center
IMS – International Monitoring System
ISC – International Seismological Center
LLNL – Lawrence Livermore National Laboratory
REB – Reviewed Event Bulletin
RSTT – Regional Seismic Travel Times
SNR – Signal-to-Noise Ratio

DISTRIBUTION LIST

DTIC/OCP
8725 John J. Kingman Rd, Suite 0944
Ft Belvoir, VA 22060-6218 1 cy

AFRL/RVIL
Kirtland AFB, NM 87117-5776 2 cys

Official Record Copy
AFRL/RVBYE/Robert Raistrick 1 cy